arXiv:1410.8816v5 [cs.CC] 19 Jan 2016

Affine reductions for LPs and SDPs

Gábor Braun[1], Sebastian Pokutta[2], and Daniel Zink[3]

[1]ISyE, Georgia Institute of Technology, Atlanta, GA, USA. *Email:* gabor.braun@isye.gatech.edu
[2]ISyE, Georgia Institute of Technology, Atlanta, GA, USA. *Email:* sebastian.pokutta@isye.gatech.edu
[3]ISyE, Georgia Institute of Technology, Atlanta, GA, USA. *Email:* daniel.zink@gatech.edu

January 20, 2016

Abstract

We define a reduction mechanism for LP and SDP formulations that degrades approximation factors in a controlled fashion. Our reduction mechanism is a minor restriction of classical reductions establishing inapproximability in the context of PCP theorems. As a consequence we establish strong linear programming inapproximability (for LPs with a polynomial number of constraints) for many problems. In particular we obtain a $\frac{3}{2} - \varepsilon$ inapproximability for VertexCover answering an open question in Chan et al. [2013] and we answer a weak version of our sparse graph conjecture posed in Braun et al. [2014a] showing an inapproximability factor of $\frac{1}{2} + \varepsilon$ for bounded degree IndependentSet. In the case of SDPs, we obtain inapproximability results for these problems relative to the SDP-inapproximability of MaxCUT. Moreover, using our reduction framework we are able to reproduce various results for CSPs from Chan et al. [2013] via simple reductions from Max-2-XOR.

1 Introduction

Linear Programming (LP) and Semidefinite Programming (SDP) formulations are ubiquitous in combinatorial optimization problems and are often an important ingredient in the construction of approximation algorithms. Extended formulations study the size of such LP or SDP formulations. So far most strong lower bounds have been obtained by ad-hoc analysis, which is stark contrast to the computational complexity world, where we often resort to reduction mechanisms to establish hardness or hardness of approximation.

In this work we establish a strong reduction mechanism for approximate LP and SDP formulations. Having a reduction mechanism in place has three main appeals.

(i) The reduction mechanism allows for the propagation of hardness results across optimization problems very similar in spirit to the computational complexity approach. In particular, no specific knowledge of how the hardness of the base problem has been established is required. As such it turns establishing LP/SDP hardness into a routine task for many problems allowing a much broader community to benefit from results from extended formulations.

(ii) Any future improvements to the strength of the lower bounds of the base problems immediately propagate to all problems that one can reduce to. In particular, it is likely that the lower bounds in Chan et al. [2013] and Lee et al. [2014a], which form the base hard problems that we reduce from can be further improved.

(iii) Having a reduction mechanism in place provides an ordering on the hardness of problems. As such it is a first step in identifying the LP/SDP analog to a complete problem. This is a quite appealing

problem in its own right as it will likely have to reconcile the LP-hardness of matching and 3-SAT, which is in contrast to the polynomial time solvability of matching in computational complexity.

Our reduction mechanism is compatible with numerous known reductions in the literature that have been used to analyze computational complexity, so that we can reuse these reductions in the context of LPs and SDPs. As the approach is very similar in all cases, i.e., to verify that the additional conditions of our mechanism are met, we will only state a select few that are of specific interest. We obtain new LP inapproximability results for problems such as, e.g., VertexCover, Max-MULTI-k-CUT, and bounded degree IndependentSet which are not 0/1 CSPs and hence not captured by the approach in Chan et al. [2013]. Moreover we reproduce previous results for CSPs in Chan et al. [2013] via direct reductions from Max-2-XOR. This is interesting in view of Max-2-XOR being the actual driver of complexity for most of these problems. In particular, establishing a stronger lower bound for Max-2-XOR, e.g., of the form $2^{\Omega(n^\delta)}$ would imply improve the LP-hardness of approximation of many other CSPs.

Our reduction mechanism is based on a more abstract view of extended formulations that is motivated by the earlier approach in Chan et al. [2013] to capture linear programming formulations independent of the specific linear encoding in the context of CSPs, where the feasible region is the whole 0/1 cube. We generalize this approach to studying arbitrary combinatorial optimization problems and the resulting model captures previous approaches studying approximations via polyhedral pairs (see Braun et al. [2012, 2014c]). In the same spirit as Chan et al. [2013] we do not rely on lifting a specific polyhedral representation but rather work directly with the combinatorial optimization problem, i.e., we answer questions of the form

Given a combinatorial optimization problem, what is the smallest size of any of its LPs or SDPs?

While this difference to the traditional extended formulation model is more of a philosophical nature and the results are equivalent to those obtained via the traditional extended formulations setup, on a technical level this perspective significantly simplifies the treatment of approximate LP/SDP formulations and it enables the formulation of the reduction mechanism.

Related work

For a detailed account on various lower bounds for extended formulations of combinatorial optimization problems, see e.g., Fiorini et al. [2012], Rothvoß [2014]. A reduction mechanism for *exact* extended formulations has been considered in Avis and Tiwary [2013], Pokutta and Van Vyve [2013], where it already provided lower bounds on the exact extension complexity of various polytopes, however both fall short to capture approximations as they cannot incorporate the necessary affine shifts. Our generalization of (encoding dependent) extended formulations to encoding independent *formulation complexity* is a natural extension of Chan et al. [2013] and Braun et al. [2014b]. The former studied an encoding-independent model for *uniform formulations of CSPs*, whereas the latter used a restricted LP version of the model presented here. Recently, there has been also progress in relating general LP and SDP formulations with hierarchies (Chan et al. [2013], Lee et al. [2014b,a]), and we reuse some of the results as the basis for later reductions. In particular Lee et al. [2014a] recently established super-polynomial lower bounds for the size of approximate SDP formulations capturing Max-3-SAT and other problems, however the SDP hardness of approximation of our base problems is still open. Our approach is also related to Kaibel et al. [2013] as well as inapproximability reductions in the context of PCPs and Lasserre hierarchies (see e.g., Håstad [1999], Trevisan et al. [2000], Håstad [2001], Trevisan [2004], Schoenebeck [2008], Tulsiani [2009]).

Contribution

Our contribution can be broadly separated into the following three parts, where the reduction mechanism is the main contribution. The abstract view on extended formulations should be considered an enabler for

the reduction mechanism and the analysis of approximate LP/SDP formulations. The use of the reduction mechanism is exemplified via various new inapproximability results.

We stress that all results are *independent* of P vs. NP and pertain to solving combinatorial problems with LPs or SDPs.

(i) **Factorization Theorem for Combinatorial Problems.** The key element in the analysis of extended formulations is Yannakakis's celebrated Factorization Theorem (see Yannakakis [1991, 1988]) and its generalizations (see e.g., Gouveia et al. [2013], Braun et al. [2012, 2014c], Chan et al. [2013]) equating the minimal size of an extended formulation with a property of a slack matrix, e.g., in the linear case the nonnegative rank. We provide an abstract, unified version (Theorem 3.5) of these factorization theorems for combinatorial optimization problems *and their approximations* acting directly on the problem. From an optimal factorization one can explicitly reconstruct an optimal encoding as a linear program or semidefinite program. As a consequence we characterize the linear programming complexity as well as the semidefinite programming complexity of a combinatorial problem *independent* of the linear representation of the problem. Our employed model is essentially equivalent to the extended formulations approach as we will see (see Section 2.1) and its main purpose is to facilitate the definition (and use) of the reduction mechanism.

(ii) **Reduction mechanism.** We provide a purely combinatorial and conceptually simple framework for reductions (similar to L-reductions) of optimization problems in the context of LPs and SDPs (and in fact *any* other conic programming paradigm), where approximations are inherent without the need of any polyhedra (and polyhedral pairs) as compared to e.g., Braun et al. [2012, 2014c]. Thereby we overcome many technical difficulties that prevented direct reductions for approximate LPs or SDPs in the past. Many reductions in the context of PCP inapproximability are compatible with our mechanism and hence can be reused. Contrasting the above, so far LP inapproximability results have been only obtained for very restricted classes of problems (see Braun et al. [2012], Chan et al. [2013], Braun et al. [2014c]) requiring a case-by-case analysis.

(iii) **LP inapproximability and conditional SDP inapproximability of specific problems.** Our reduction mechanism opens up the possibility to reuse previous hardness results to establish inapproximability of problems. As a case in point, we establish the first LP inapproximability result for VertexCover and IndependentSet as well as reproduce the results in Chan et al. [2013], all by simple and direct reductions from the LP-inapproximability of Max-2-XOR within a factor better than $\frac{1}{2}$ established in Chan et al. [2013]. For SDP formulations our reductions establish relative inapproximability between the considered problems as strong SDP hardness of approximation for our base problems Max-k-XOR are unknown. Several of our inapproximability results are new and others reproduce the results in Chan et al. [2013] via direct reductions from a single hard problem. The conditional SDP inapproximability factors are formulated under the assumption that the Goemans-Williams SDP for MaxCUT is optimal (see Conjecture 4.6), which is compatible with the Unique Games Conjecture. Alternatively, these reductions can be also combined with the $15/16 + \varepsilon$ SDP-hardness of approximation for MaxCUT that was recently established in Braun et al. [2015c] to obtain unconditional (but weaker) inapproximability factors. To obtain the respective factors it suffices to replace c_{GW} with $15/16$ in the arguments.

In particular, we answer an open question regarding the inapproximability of VertexCover (see Chan et al. [2013]) and we answer a weak version of our sparse graph conjecture posed in Braun et al. [2014a]. We obtain results as provided in the following table:

Problem	Inapproximability (LP)	Inapproximability (SDP under 4.6)	Inapproximability (PCP)	Approximability (LP)
VertexCover (new)	$\frac{3}{2}-\varepsilon$	$1.12144-\varepsilon$	$1.361-\varepsilon$	2
Min-2-CNFDeletion (new)	$\omega(1)$	$\omega(1)$	$2.889-\varepsilon$	—
MinUnCUT (new)	$\omega(1)$	$\omega(1)$	$C_{\text{uncut}}-\varepsilon$	—
Max-MULTI-k-CUT (new)	$\frac{2c(k)+1}{2c(k)+2}+\varepsilon$	$\frac{c(k)+c_{GW}}{c(k)+1}$	$1-\frac{1}{34k}+\varepsilon$	$\frac{1}{2(1-1/k)}$
bounded degree IndependentSet (new)	$\frac{1}{2}+\varepsilon$	$c_{GW}+\varepsilon$	$O\left(\frac{\log^4\Delta}{\Delta}\right)$	—
Max-2-SAT	$\frac{3}{4}+\varepsilon$ (optimal)	$\frac{1+c_{GW}}{2}+\varepsilon\approx 0.93928+\varepsilon$	$\frac{21}{22}+\varepsilon$	$\frac{3}{4}$
Max-3-SAT	$\frac{3}{4}+\varepsilon$	$\frac{1+c_{GW}}{2}+\varepsilon\approx 0.93928+\varepsilon$	$\frac{7}{8}+\varepsilon$	$\frac{19}{27}$
Max-DICUT	$\frac{1}{2}+\varepsilon$ (optimal)	$c_{GW}+\varepsilon$	$\frac{12}{13}+\varepsilon$	$\frac{1}{2}$
Max-2-CONJSAT	$\frac{1}{2}+\varepsilon$ (optimal)	$c_{GW}+\varepsilon$	$\frac{9}{10}+\varepsilon$	$\frac{1}{2}$

Here $c_{GW}\approx 0.87856$ is the approximation factor of the algorithm for MaxCUT from Goemans and Williamson [1995], and $c(k)$ is a constant depending on k defined in Section 6.1.

Note that our conditional SDP inapproximability of Max-2-SAT and Max-3-SAT within is $(1+c_{GW})/2+\varepsilon\approx 0.93928+\varepsilon$, better than the current best unconditional inapproximability of $7/8$ from [Lee et al., 2014a, Theorem 1.5]; see Section 6 for details. Moreover, there exists a linear program for IndependentSet achieving an approximation guarantee of $2\sqrt{n}$, which is indeed better than the $n^{1-\varepsilon}$ hardness of approximation by Håstad [1999], so that we cannot expect to obtain a $n^{1-\varepsilon}$ hardness of approximation ; see Bazzi et al. [2015] for a detailed discussion.

As a nice (minor) byproduct, the presented framework augments the results in Fiorini et al. [2012], Braun et al. [2012], Braverman and Moitra [2013], Braun et al. [2014c], Rothvoß [2014] to be independent of the chosen linear encoding (see Section 3.1), which allows us to reuse these results directly in reductions. Also, approximations of the slack matrix provide approximate programs for the optimization problem, as shown in Theorem 7.2. Our approach also immediately carries over to symmetric formulations (see Braun et al. [2015a] for SDPs) and other conic programming paradigms; the details are left to the interested reader.

Finally, we would also like to note that our model and reductions have been already applied to obtain inapproximability results for approximate GraphIsomorphism in Braun et al. [2015b], exponential lower bound for symmetric SDPs for Matching (see [Braun et al., 2015a, Theorem 3.1]), and$(2-\varepsilon)$-inapproximability for VertexCover together with inapproximability of IndependentSet within any constant factor in Bazzi et al. [2015] improving our Theorem 5.3 by a significantly more involved argument. An extension of this reduction mechanism is presented in Braun et al. [2015c] relaxing some of our conditions, which enables the study of fractional optimization problems (such as e.g., Sparsest Cut).

2 Optimization problems

We intend to study the required size of a linear program or semidefinite program capturing a combinatorial optimization problem with specified approximation guarantees. In our context an optimization problem is defined as follows.

Definition 2.1 (Optimization problems)**.** An *optimization problem* $\mathcal{P} = (\mathcal{S}, \mathcal{F}, \mathrm{val})$ consists of a set $\mathcal{S}$ of *feasible solutions* and a set $\mathcal{F}$ of *instances*, together with a real-valued objective function $\mathrm{val}\colon \mathcal{S} \times \mathcal{F} \to \mathbb{R}$.

A wide class of examples consist of constraint satisfaction problems (CSPs):

Definition 2.2 (Maximum Constraint Satisfaction Problem (CSP))**.** A *constraint family* $\mathcal{C} = \{C_1, \ldots, C_m\}$ on the boolean variables $x_1, \ldots, x_n$ is a family of boolean functions C_i in $x_1, \ldots, x_n$. The C_i are *constraints* or *clauses*. The problem $\mathcal{P}(\mathcal{C})$ corresponding to a constraint family $\mathcal{C}$ has

(i) **feasible solutions** all 0/1 assignments s to $x_1, \ldots, x_n$;

(ii) **instances** all nonnegative weightings $w_1, \ldots, w_m$ of the constraints $C_1, \ldots, C_m$

(iii) **objective function** the weighted sum of satisfied constraints: $\mathrm{sat}_{w_1,\ldots,w_m}(s) = \sum_i w_i C_i(s)$.

The goal is to maximize the weights of satisfied constraints, in particular CSPs are maximization problems. A *maximum Constraint Satisfaction Problem* is an optimization problem $\mathcal{P}(\mathcal{C})$ for some constraint family $\mathcal{C}$. A k-CSP is a CSP where every constraint depends on at most k variables.

For brevity, we shall simply use CSP for a maximum CSP, when there is no danger of confusion with a minimum CSP. In the following, we shall restrict to instances with 0/1 weights, i.e., an instance is a subset $L \subseteq \mathcal{C}$ of constraints, and the objective function computes the number $\mathrm{sat}_L(s) = \sum_{C \in L} C(s)$ of constraints in L satisfied by assignment s. Restriction to specific instances clearly does not increase formulation complexity.

As a special case, the Max-k-XOR problem restricts to constraints, which are XORs of k literals. Here we shall write the constraints in the equivalent equation form $x_{i_1} \oplus \cdots \oplus x_{i_k} = b$, where $\oplus$ denotes the addition modulo 2.

Definition 2.3 (Max-k-XOR)**.** For fixed k and n, the problem Max-k-XOR is the CSP for variables $x_1, \ldots, x_n$ and the family $\mathcal{C}$ of all constraints of the form $x_{i_1} \oplus \cdots \oplus x_{i_k} = b$ with $1 \leq i_1 < \cdots < i_k \leq n$ and $b \in \{0, 1\}$.

An even stronger important restriction is MaxCUT, a subproblem of Max-2-XOR as we will see soon. The aim is to determine the maximum size of cuts for all graphs G with $V(G) = [n]$.

Definition 2.4 (MaxCUT)**.** The problem MaxCUT has instances all simple graphs G with vertex set $V(G) = [n]$, and feasible solutions all cuts on $[n]$, i.e., functions $s\colon [n] \to \{0, 1\}$. The objective function val computes the number of edges $\{i, j\}$ of G cut by the cut, i.e., with $s(i) \neq s(j)$.

The problem MaxCUT_Δ is the subproblem of MaxCUT considering only graphs G with maximum degree at most Δ.

We have $\mathrm{val}_G(s) = \mathrm{sat}_{L(G)}(s)$, for the constraint set $L(G) = \left\{ x_i \oplus x_j = 1 \,\middle|\, \{i, j\} \in E(G) \right\}$, realizing MaxCUT as a subproblem of Max-2-XOR with the same feasible solutions.

We are interested in approximately solving an optimization problem $\mathcal{P}$ by means of a linear program or a semidefinite program. Recall that a typical PCP inapproximability result states that it is hard to decide between $\max \mathrm{val}_i \leq S(i)$ and $\max \mathrm{val}_i \geq C(i)$ for a class of instances i and some easy-to compute functions S and C usually refereed to as *soundness* and *completeness*. Here and below $\max \mathrm{val}_i$ denotes the maximum value of the function val_i over the respective set of feasible solutions. We adopt the terminology to linear programs and semidefinite programs. We start with the linear case.

Definition 2.5 (LP formulation of an optimization problem)**.** Let $\mathcal{P} = (\mathcal{S}, \mathcal{F}, \mathrm{val})$ be an optimization problem with real-valued functions C, S on $\mathcal{F}$, called *completeness guarantee* and *soundness guarantee*, respectively. If $\mathcal{P}$ is a maximization problem, then let $\mathcal{F}^S := \left\{ f \in \mathcal{F} \,\middle|\, \max \mathrm{val}_f \leq S(f) \right\}$ denote the set of

instances, for which the maximum is upper bounded by soundness guarantee S. If $\mathcal{P}$ is a minimization problem, then let $\mathcal{F}^S := \left\{ f \in \mathcal{F} \,\middle|\, \min \mathrm{val}_f \geq S(f) \right\}$ denote the set of instances, for which the minimum is lower bounded by soundness guarantee S.

A (C, S)-*approximate LP formulation* of $\mathcal{P}$ is a linear program $Ax \leq b$ with $x \in \mathbb{R}^d$ together with the following *realizations*:

(i) **Feasible solutions** as vectors $x^s \in \mathbb{R}^d$ for every $s \in \mathcal{S}$ so that

$$Ax^s \leq b \qquad \text{for all } s \in \mathcal{S}, \tag{1}$$

i.e., the system $Ax \leq b$ is a relaxation (superset) of $\mathrm{conv}\left(x^s \mid s \in \mathcal{S}\right)$.

(ii) **Instances** as affine functions $w^f : \mathbb{R}^d \to \mathbb{R}$ for every $f \in \mathcal{F}^S$ such that

$$w^f(x^s) = \mathrm{val}_f(s) \qquad \text{for all } s \in \mathcal{S}, \tag{2}$$

i.e., we require that the linearization w^f of val_f is exact on all x^s with $s \in \mathcal{S}$.

(iii) **Achieving guarantee** C via requiring

$$\max \left\{ w^f(x) \,\middle|\, Ax \leq b \right\} \leq C(f) \qquad \text{for all } f \in \mathcal{F}^S, \tag{3}$$

for maximization problems (resp. $\min \left\{ w^f(x) \,\middle|\, Ax \leq b \right\} \geq C(f)$ for minimization problems).

The *size* of the formulation is the number of inequalities in $Ax \leq b$. Finally, the *LP formulation complexity* $\mathrm{fc}_+(\mathcal{P}, C, S)$ of the problem $\mathcal{P}$ is the minimal size of all its LP formulations.

For all instances $f \in \mathcal{F}$ soundness and completeness should satisfy $C(f) \geq S(f)$ in the case of maximization problems and $C(f) \leq S(f)$ in the case of minimization problems in order to capture the notion of relaxations and we assume this condition in the remainder of the paper.

Remark 2.6. We use affine maps instead of linear maps to allow easy shifting of functions. At the cost of an extra dimension and an extra equation, affine functions can be realized as linear functions.

Remark 2.7 (Inequalities vs. Equations). Traditionally in extended formulations, one would separate the description into equations and inequalities and one would only count inequalities. In our framework, equations can be eliminated by restricting to the affine space defined by them, and parametrizing it as a vector space. However, note that restricting to linear functions, one might need an equation to represent affine functions by linear functions.

For determining the exact maximum of a maximization problem, one chooses $C(f) = S(f) := \max \mathrm{val}_f$. To show inapproximability within an approximation factor $0 < \rho \leq 1$, one chooses guarantees satisfying $\rho C(f) \geq S(f)$. This choice is motivated to be comparable with factors of approximation algorithms finding a feasible solution s with $\mathrm{val}_f(s) \geq \rho \max \mathrm{val}_f$. For minimization problem, $C(f) = S(f) := \min \mathrm{val}_f$ in the exact case, and $\rho C(f) \leq S(f)$ for an approximation factor $\rho \geq 1$ provided val_f is nonnegative. This model of *outer* approximation streamlines the models in Chan et al. [2013], Braun et al. [2014b], Lee et al. [2014b], and also captures, simplifies, and generalizes approximate extended formulations from Braun et al. [2012, 2014c]; see Section 2.1 for a discussion.

We will now adjust Definition 2.5 to the semidefinite case. For symmetric matrices, as customary, we use the Frobenius product as scalar product, i.e., $\langle A, B \rangle = \mathrm{Tr}[AB]$. Recall that the psd-cone is self-dual under this scalar product.

Definition 2.8 (SDP formulation of an optimization problem)**.** Let $\mathcal{P} = (\mathcal{S}, \mathcal{F}, \mathrm{val})$ be a maximization problem with real-valued functions C, S on $\mathcal{F}$ and let $\mathcal{F}^{\mathcal{S}} := \left\{ f \in \mathcal{F} \,\middle|\, \max \mathrm{val}_f \leq S(f) \right\}$ as in Definition 2.5.

A (C, S)-approximate *SDP formulation* of $\mathcal{P}$ consists of a linear map $\mathcal{A} \colon \mathbb{S}^d \to \mathbb{R}^k$ and a vector $b \in \mathbb{R}^k$ (defining a semidefinite program $\left\{ X \in \mathbb{S}^d_+ \,\middle|\, \mathcal{A}(X) = b \right\}$). Moreover, we require the following *realizations* of the components of $\mathcal{P}$:

(i) **Feasible solutions** as vectors $X^s \in \mathbb{S}^d_+$ for every $s \in \mathcal{S}$ so that

$$\mathcal{A}(X^s) = b \tag{4}$$

i.e., the system $\mathcal{A}(X) = b, X \in \mathbb{S}^d_+$ is a relaxation of $\mathrm{conv}\,(X^s \,|\, s \in \mathcal{S})$.

(ii) **Instances** as affine functions $w^f \colon \mathbb{S}^d \to \mathbb{R}$ for every $f \in \mathcal{F}^{\mathcal{S}}$ with

$$w^f(X^s) = \mathrm{val}_f(s) \qquad \text{for all } s \in \mathcal{S}, \tag{5}$$

i.e., we require that the linearization w^f of val_f is exact on all X^s with $s \in \mathcal{S}$.

(iii) **Achieving guarantee** C via requiring

$$\max \left\{ w^f(X) \,\middle|\, \mathcal{A}(X^s) = b,\ X^s \in \mathbb{S}^d_+ \right\} \leq C(f) \qquad \text{for all } f \in \mathcal{F}, \tag{6}$$

for maximization problems, and the analogous inequality for minimization problems.

The *size* of the formulation is the parameter d. The *SDP formulation complexity* $\mathrm{fc}_{\oplus}(\mathcal{P}, C, S)$ of the problem $\mathcal{P}$ is the minimal size of all its SDP formulations.

2.1 Relation to approximate extended formulations

Traditionally in extended formulations, one would start from an initial polyhedral representation of the problem and bound the size of its smallest possible lift in higher-dimensional space. In the linear case for example, the minimal number of required inequalities would constitute the extension complexity of that polyhedral representation. Our notion of *formulation complexity* can be understood as the minimum extension complexity over all possible polyhedral encodings of the optimization problem. This independence of encoding addresses previous concerns that the obtained lower bounds are polytope-specific or encoding-specific and alternative linear encodings (i.e., different initial polyhedron) of the same problem might admit smaller formulations: we showed that this is not the case. More precisely, in view of the results from above the standard notion of extension complexity and formulation complexity are essentially equivalent, however the more abstract perspective simplifies the handling of approximations and reductions as we will see in Section 4.

The notion of *LP formulation*, its size, and LP formulation complexity are closely related to *polyhedral pairs* and linear encodings (see Braun et al. [2012, 2014c], and also Pashkovich [2012]). In particular, given a (C, S)-approximate LP formulation of a maximization problem $\mathcal{P}$ with linear program $Ax \leq b$, representations $\left\{ x^s \,|\, s \in \mathcal{S} \right\}$ of feasible solutions and $\left\{ w^f \,\middle|\, f \in \mathcal{F}^{\mathcal{S}} \right\}$ of instances, one can define a polyhedral pair encoding $\mathcal{P}$

$$P := \mathrm{conv}\,(x^s \,|\, s \in \mathcal{S}),$$
$$Q := \left\{ x \in \mathbb{R}^d \,\middle|\, \langle w^f, x \rangle \leqslant C(f),\ \forall f \in \mathcal{F}^{\mathcal{S}} \right\}.$$

Then for $K := \{x \,|\, Ax \leq b\}$, we have $P \subseteq K \subseteq Q$. Note that there is no need for the approximating polyhedron K to reside in extended space, as P and Q already live there.

Put differently, the LP formulation complexity of $\mathcal{P}$ is the minimum size of an extended formulation over all possible linear encodings of $\mathcal{P}$. The semidefinite case is similar, with the only difference being that K is now a spectrahedron, being represented by a semidefinite program instead of a linear program.

3 Factorization theorem and slack matrix

We provide an algebraic characterization of formulation complexity via the *slack matrix of an optimization problem*, similar in spirit to factorization theorems for extended formulations (see e.g., Yannakakis [1991, 1988], Gouveia et al. [2013], Braun et al. [2012, 2014c]), with a fundamental difference pioneered in Chan et al. [2013] that there is no linear system to start from. The linear or semidefinite program is constructed from scratch using a matrix factorization. This also extends Braun et al. [2014b], by allowing affine functions, and using a modification of nonnegative rank, to show that formulation complexity depends only on the slack matrix.

Definition 3.1 (Slack matrix of $\mathcal{P}$)**.** Let $\mathcal{P} = (\mathcal{S}, \mathcal{F}, \mathrm{val})$ be an optimization problem with guarantees C, S. The (C, S)*-approximate slack matrix* of $\mathcal{P}$ is the nonnegative $\mathcal{F}^{\mathcal{S}} \times \mathcal{S}$ matrix M, with entries

$$M(f,s) \coloneqq \begin{cases} C(f) - \mathrm{val}_f(s) & \text{if } \mathcal{P} \text{ is a maximization problem,} \\ \mathrm{val}_f(s) - C(f) & \text{if } \mathcal{P} \text{ is a minimization problem.} \end{cases}$$

We introduce the *LP factorization* of a nonnegative matrix, which for slack matrices captures the LP formulation complexity of the underlying problem.

Definition 3.2 (LP factorization of a matrix)**.** A *size-r LP factorization* of $M \in \mathbb{R}_+^{m \times n}$ is a factorization $M = TU + \mu\mathbb{1}$ where $T \in \mathbb{R}_+^{m \times r}$, $U \in \mathbb{R}_+^{r \times n}$ and $\mu \in \mathbb{R}_+^{m \times 1}$. Here $\mathbb{1}$ is the $1 \times n$ matrix with all entries being 1. The *LP rank* $\mathrm{rank}_{\mathrm{LP}} M$ *of* M is the minimum r such that there exists a size-r LP factorization of M.

A size-r LP factorization is equivalent to a decomposition $M = \sum_{i \in [r]} u_i v_i^{\mathsf{T}} + \mu\mathbb{1}$ for some (column) vectors $u_i \in \mathbb{R}_+^m$, $v_i \in \mathbb{R}_+^n$ with $i \in [r]$ and a column vector $\mu \in \mathbb{R}_+^m$. It is a slight modification of a nonnegative matrix factorization, disregarding simultaneous shift of all columns by the same vector, i.e., allowing an additional term $\mu \cdot \mathbb{1}$ *not* contributing to the size, so clearly, $\mathrm{rank}_{\mathrm{LP}} \leq \mathrm{rank}_+ M \leq \mathrm{rank}_{\mathrm{LP}} M + 1$.

One similarly defines SDP factorizations of nonnegative matrices.

Definition 3.3. A *size-r SDP factorization* of $M \in \mathbb{R}_+^{m \times n}$ is a factorization is a collection of matrices $T_1, \ldots, T_m \in \mathbb{S}_+^r$ and $U_1, \ldots, U_n \in \mathbb{S}_+^r$ together with $\mu \in \mathbb{R}_+^{m \times 1}$ so that $M_{ij} = \mathrm{Tr}[T_i U_j] + \mu(i)$. The *SDP rank* $\mathrm{rank}_{\oplus} M$ *of* M is the minimum r such that there exists a size-r SDP factorization of M.

For the next theorem, we need the folklore formulation of linear duality using affine functions, see e.g., [Schrijver, 1986, Corollary 7.1h].

Lemma 3.4 (Affine form of Farkas's Lemma)**.** *Let $P \coloneqq \{x \mid A_j x \leq b_j, j \in [r]\}$ be a non-empty polyhedron. An affine function Φ is nonnegative on P if and only if there are nonnegative multipliers λ_j, λ_0 with*

$$\Phi(x) \equiv \lambda_0 + \sum_{j \in [r]} \lambda_j (b_j - A_j x).$$

We are ready for the factorization theorem for optimization problems.

Theorem 3.5 (Factorization theorem for formulation complexity)**.** *Consider an optimization problem $\mathcal{P} = (\mathcal{S}, \mathcal{F}, \mathrm{val})$ with (C, S)-approximate slack matrix M. Then we have*

$$\mathrm{fc}_+(\mathcal{P}, C, S) = \mathrm{rank}_{\mathrm{LP}} M, \qquad \textit{and} \qquad \mathrm{fc}_{\oplus}(\mathcal{P}, C, S) = \mathrm{rank}_{\mathrm{SDP}} M.$$

for linear formulations and semidefinite formulations, respectively.

Remark 3.6. The factorization theorem for polyhedral pairs (see Braun et al. [2012], Pashkovich [2012], Braun et al. [2014c]) states that the nonnegative rank and extension complexity might differ by 1, which was slightly elusive. Theorem 3.5 clarifies, that this is the difference between the LP rank and nonnegative rank, i.e., formulation complexity is a *property of the slack matrix*. Note also that for the slack matrix of a polytope, every row contains a 0 entry, and hence the $\mu\mathbb{1}$ term in any LP factorization must be 0. Therefore the nonnegative rank and LP rank coincide for polytopes. Similar remarks apply to SDP factorizations.

Proof of Theorem 3.5—the linear case. We will confine ourselves to the case of $\mathcal{P}$ being a maximization problem. For minimization problems, the proof is analogous.

To prove $\operatorname{rank}_{\mathrm{LP}} M \leq \operatorname{fc}_+(\mathcal{P}, C, S)$, let $Ax \leq b$ be an arbitrary (C,S)-approximate, size-r LP formulation of $\mathcal{P}$, with realizations $\left\{w^f \,\middle|\, f \in \mathcal{F}^{\mathcal{S}}\right\}$ of instances and $\{x^s \mid s \in \mathcal{S}\}$ of feasible solutions. We shall construct a size-r nonnegative factorization of M. As $\max_{x:Ax\leq b} w^f(x) \leq C(f)$ by Condition (3), via the affine form of Farkas's lemma, Lemma 3.4 we have

$$C(f) - w^f(x) = \sum_{j=1}^{r} T(f,j)\left(b_j - \langle A_j, x\rangle\right) + \mu(f)$$

for some nonnegative multipliers $T(f,j), \mu(f) \in \mathbb{R}_+$ with $1 \leq j \leq r$. By taking $x = x^s$, we obtain

$$M(f,s) = \sum_{j=1}^{r} T(f,j)U(j,s) + \mu(f), \qquad \text{with} \quad U(j,s) \coloneqq b_j - \langle A_j, x^s\rangle \qquad \text{for } j > 0.$$

i.e., $M = TU + \mu\mathbb{1}$. By construction, T and μ are nonnegative. By Condition (1) we also obtain that U is nonnegative. Therefore $M = TU + \mu\mathbb{1}$ is a size-r LP factorization of M.

For the converse, i.e., $\operatorname{rank}_{\mathrm{LP}} \geq \operatorname{fc}_+(\mathcal{P}, C, S)$, let $M = TU + \mu\mathbb{1}$ be a size-r LP factorization. We shall construct an LP formulation of size r. Let T_f denote the f-row of T for $f \in \mathcal{F}^{\mathcal{S}}$, and U_s denote the s-column of U for $s \in \mathcal{S}$. We claim that the linear system $x \geq 0$ with representations

$$w^f(x) \coloneqq C(f) - \mu(f) - T_f x \qquad \forall f \in \mathcal{F}^{\mathcal{S}} \qquad \text{and} \qquad x^s \coloneqq U_s \qquad \forall s \in \mathcal{S}$$

satisfies the requirements of Definition 2.5. Condition (2) is implied by the factorization $M = TU + \mu\mathbb{1}$:

$$w^f(x^s) = C(f) - \mu(f) - T_f U_s = C(f) - M(f,s) = \operatorname{val}_f(s).$$

Moreover, $x^s \geq 0$, because U is nonnegative, so that Condition (1) is fulfilled. Finally, Condition (3) also follows readily:

$$\max\left\{w^f(x) \,\middle|\, x \geq 0\right\} = \max\left\{C(f) - \mu(f) - T_f x \,\middle|\, x \geq 0\right\} = C(f) - \mu(f) \leq C(f),$$

as the nonnegativity of T implies $T_f x \geq 0$; equality holds e.g., for $x = 0$. Recall also that $\mu(f) \geq 0$. Thus we have constructed an LP formulation with r inequalities, as claimed. □

Remark 3.7. It is counter-intuitive that 0 is always a maximizer, and, actually, it is an artifact of the construction. At a conceptual level, the polyhedron $Ax \leq b$ containing $\operatorname{conv}(x^s \mid s \in \mathcal{S})$ is represented as the intersection of the nonnegative cone with an affine subspace in the slack space. The affine functions w^f are extended to attain their optimum value on this intersection in the nonnegative cone, and thus also at 0, the apex of the cone. In particular, intersecting with the affine subspace is no longer needed. See Figure 1 for an illustration.

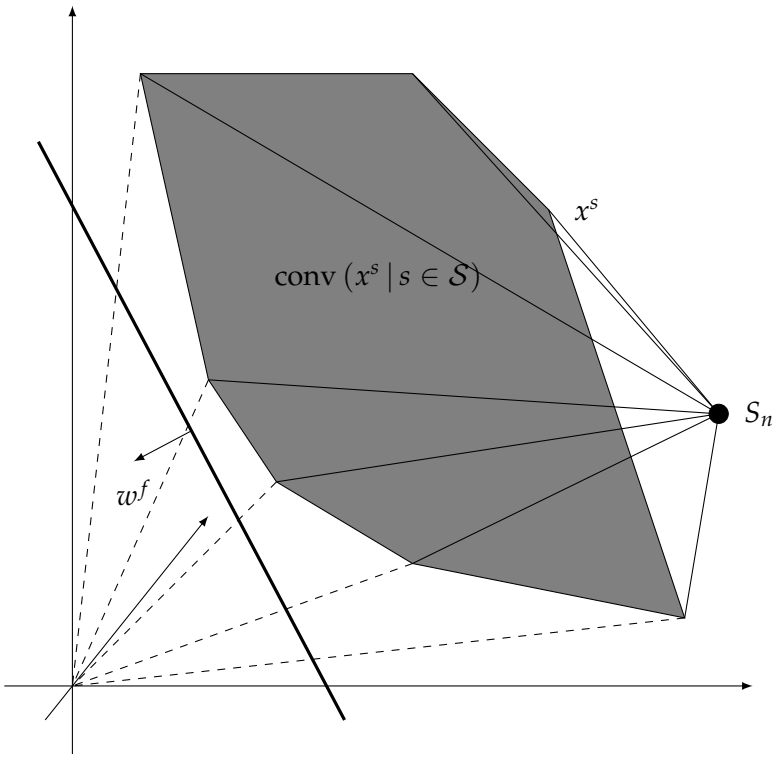

Figure 1: Linear program from an LP factorization. The LP is the positive orthant $x \geq 0$. The point 0 is a maximizer for all linearizations $w^f = C(f) - \mu(f) - T_f x$ of objective functions val_f for all instances f. The normals of all objective functions point in nonpositive direction as $T_f \geq 0$.

Remark 3.8 (Solution structure). Observe that the obtained LP formulation via the LP-factorization of the slack matrix also separates solutions $s \in S$ into two disjoint classes $S = S_o \uplus S_n$, where S_o contains those solutions that potentially can be optimal for some function, i.e., it is the set of coordinate-wise minimal points in S. The set S_n is the set of solution that are *never* optimal for any $f \in F$ as they are coordinate-wise dominated by at least one point in S_o, i.e., for any $s_n \in S_n$ there exists $s_o \in S_o$ with $s_o \leq s_n$ coordinate-wise.

This might have applications in the context of inverse optimization (see e.g., Ahuja and Orlin [2001]), where we would like to decide whether for a given solution $s \in \mathcal{S}$, there exists an $f \in \mathcal{F}$, that is maximized at s. This can now be read off the factorization.

Proof of Theorem 3.5—the semidefinite case. As before we confine ourselves to the case of $\mathcal{P}$ being a maximization problem. The proof is analogous to the linear case, but for the sake of completeness, we provide a full proof.

To prove $\mathrm{rank}_{\mathrm{SDP}} M \leq \mathrm{fc}_{\oplus}(\mathcal{P}, C, S)$, let $\mathcal{A}(X) = b, X \in \mathbb{S}^r_+$ be an arbitrary size-r SDP formulation of $\mathcal{P}$, with realizations $\left\{w^f \;\middle|\; f \in \mathcal{F}^{\mathcal{S}}\right\}$ of instances and $\{X^s \mid s \in \mathcal{S}\}$ of feasible solutions. To apply strong duality, we may assume that the convex set $\{X \in \mathbb{S}^r_+ \mid \mathcal{A}(X) = b\}$ has an interior point because otherwise it would be contained in a proper face of $\mathbb{S}^r_+$, which is an SDP cone of smaller size. We shall construct a size-r SDP factorization of M. As $\max_{X \in \mathbb{S}^r_+ \colon \mathcal{A}(X) = b} w^f(X) \leq C(f)$ by Condition (6), via the affine form of strong duality, we have

$$C(f) - w^f(X) = \langle T_f, X \rangle + \langle y^f, b - \mathcal{A}(X) \rangle + \lambda_f$$

for all $f \in \mathcal{F}$ with some $T_f \in \mathbb{S}^r_+$, $y^f \in \mathbb{R}^k$ and $\lambda_f \in \mathbb{R}_+$. By substituting X^s into X, we obtain

$$M(f,s) = C(f) - w^f(X^s) = \langle T_f, X^s \rangle + \langle y^f, b - \mathcal{A}(X^s) \rangle + \lambda_f = \langle T_f, X^s \rangle + \lambda_f,$$

which is an SDP factorization of size r.

For the converse, i.e., $\mathrm{fc}_\oplus(\mathcal{P}, C, S) \leq \mathrm{rank}_{\mathrm{SDP}} M$, let $M(f,s) = \langle T_f, U_s \rangle + \mu(f)$ be a size-r SDP factorization. We shall construct an SDP formulation of size r. We claim that the SDP formulation:

$$X \in \mathbb{S}^r_+ \tag{7}$$

with representations

$$w^f(X) \coloneqq C(f) - \mu(f) - \langle T_f, X \rangle \qquad \forall f \in \mathcal{F}^{\mathcal{S}} \qquad \text{and} \qquad X^s \coloneqq U_s \qquad \forall s \in \mathcal{S}$$

satisfies the requirements of Definition 2.8. Condition (5) follows by:

$$w^f(X^s) = C(f) - \mu(f) - \langle T_f, U_s \rangle = C(f) - M(f,s) = \mathrm{val}_f(s).$$

Moreover, the $X^s = U_s$ are psd, hence clearly satisfy the system (7), so that Condition (4) is fulfilled. Finally, Condition (6) also follows readily.

$$\max \left\{ w^f(X) \,\middle|\, X \in \mathbb{S}^r_+ \right\} = \max \left\{ C(f) - \mu(f) - \langle T_f, X \rangle \,\middle|\, X \in \mathbb{S}^r_+ \right\} = C(f) - \mu(f) \leq C(f),$$

as T_f and X being psd implies $\langle T_f, X \rangle \geq 0$; equality holds e.g., for $X = 0$. Thus we have constructed an SDP formulation of size r as claimed. □

3.1 Examples

In the context of optimization problems we typically differentiate two types of formulations. The *uniform model* asks for a formulation *for a whole family of instances*. Our Examples 3.9, 4.7 and 3.10 are all uniform models. The *non-uniform model* asks for a formulation for the weighted version of *a specific problem*, where the instances differ only in the weighting. Lower bounds or inapproximability factors for non-uniform models are usually stronger statements, as in the non-uniform case the formulation potentially could adapt to the instance resulting in potentially smaller formulations; see Bazzi et al. [2015] for such an example in the context of stable sets. We refer the reader to Chan et al. [2013], Braun et al. [2014b] for an in-depth discussion.

The difference between uniform and non-uniform sometimes depends on the point of view. For graph problems, often the non-uniform model for a graph G induces a uniform model for the family of all of its (induced) subgraphs by choosing 0/1 weights (see e.g., Definition 5.1). In other words, the non-uniform model is actually a uniform model for the class of all subinstances of G. Examples 3.9 and 4.7 with $0/1$ weights demonstrate this: they are uniform models for all subgraphs $G \subseteq K_n$, but can also be viewed as non-uniform models for K_n.

Note that these models can also be used for studying average case complexity. For example, one might consider the complexity of the problem for a randomly selected large class of instances using the uniform model, or one might consider the non-uniform model for a randomly selected instance. For the maximum stable set problem, both random versions were examined in Braun et al. [2014b].

3.1.1 The matching problem revisited

The lower bounds in Fiorini et al. [2012], Rothvoß [2014] are concerned with specific polytopes, namely the TSP polytope as well as the matching polytope. We obtain, as a slight generalization, the same lower bounds for the Hamiltonian cycle problem (which is captured by the TSP problem with appropriate weights) as well as the matching *problem*, independent of choice of the specific polytope that represents the encoding. Here we present only the matching problem. The lower bound for the Hamiltonian cycle problem will be obtained in Section 4.1 via a reduction.

Example 3.9 (Maximum matching problem). The maximum matching problem $\mathcal{P}_{\text{Match}}$ asks for the maximum size of matchings in a given graph. While it can be solved in polynomial time, the matching polytope has exponential extension complexity as shown in Rothvoß [2014]. Using the framework from above, we immediately obtain that the matching problem has high LP formulation complexity, reusing the lower bound on the nonnegative rank of the slack matrix of the matching polytope.

For the natural formulation of the problem in our framework, let n be fixed, and let the set $\mathcal{S}$ of feasible solutions consist of all perfect matchings M of the complete graph on $2n$ vertices, and the instances be all (simple) graphs G on $[2n]$. The value $\text{val}_G(M)$ for a graph G and a perfect matching M is defined to be

$$\text{val}_G(M) := |M \cap E(G)|$$

the number of edges shared by M and G, i.e., the size of the matching $M \cap E(G)$ of G. Clearly, all maximum matchings of G can be obtained this way (via extension to any perfect matching on $[2n]$). Thus $\max \text{val}_G$ is the matching number $\nu(G)$ of G, the size of the maximum matchings of G.

Inspired by the description of the facets of the matching polytope in Edmonds [1965], we only consider complete subgraphs on odd-sized subsets U. For such a complete graph K_U on the odd-sized set U, we have $\max \text{val}_{K_U} = \frac{|U|-1}{2}$. Let $\delta(U)$ denote the set of all edges between U and its complement $[2n] \setminus U$. We have the identity $|U| = 2\,|M \cap E(K_U)| + |M \cap \delta(U)|$ and thus obtain the slack matrix for the exact problem (i.e., $C(G) = \max \text{val}_G$)

$$S(K_U, M) := C(K_U) - \text{val}_{K_U} = \frac{|U|-1}{2} - |M \cap E(K_U)| = \frac{|M \cap \delta(U)| - 1}{2}.$$

This submatrix has nonnegative rank $2^{\Omega(n)}$ by Rothvoß [2014] and hence the LP formulation complexity of the maximum matching problem is $2^{\Omega(n)}$, i.e., $\text{fc}_+(\mathcal{P}_{\text{Match}}) = 2^{\Omega(n)}$.

The result can be extended to the approximate case with an approximation factor $(1+\varepsilon/n)^{-1}$, by invoking the lower bound for the resulting slack matrix from Braun and Pokutta [2015], showing that the maximum matching problem does not admit any fully-polynomial size relaxation scheme. Note that this is not unexpected as for the maximum matching problem an approximation factor of about $1 - \frac{\varepsilon}{n}$ corresponds to an error of less than one edge in the unweighted case for small ε, so that the decision problem could be decided via the approximation. This is a behavior similar to FPTAS and strong NP-hardness for combinatorial optimization problems that are mutually exclusive (under standard assumptions).

3.1.2 Independent set problem

We provide an example for maximum independent sets in a uniform model; see Braun et al. [2014b] for more details as well as an average case analysis. Here there is no bound on the maximum degree of graphs, unlike in Theorem 5.3.

Example 3.10 (Maximum independent set problem (uniform model)). Let us consider the maximum independent set problem $\mathcal{P}$ over some family $\mathcal{G}$ of graphs G where $V(G) \subseteq [n]$ with aim to estimate the maximum size $\alpha(G)$ of independent sets in each $G \in \mathcal{G}$.

A natural choice is to let the *feasible solutions* be all subsets S of $[n]$, and the *instances* be all $G \in \mathcal{G}$. The objective function is

$$\mathrm{val}_G(S) := |V(G) \cap S| - |E(G(S))|\,.$$

Here $\mathrm{val}_G(S)$ can be easily seen to lower bound the size of an independent set, obtained from S by removing vertices not in G, and also removing one end point of every edge with both end points in S. Clearly, $\mathrm{val}_G(S) = |S|$ for independent sets S of G, i.e., in this case our choice is exact. Thus $\alpha(G) = \max_{S\subseteq[n]} \mathrm{val}_G(S)$.

Let us consider the special case when $\mathcal{G}$ is the set of all simple graphs with $V(G) \subseteq [n]$. We shall use guarantees $S(G) := \max \mathrm{val}_G = \alpha(G)$ and $C(G) := \rho^{-1}\,\mathrm{val}_G$ for an approximation factor $0 < \rho \leq 1$. Restricting to complete graphs K_U with $U \subseteq [n]$, the obtained slack matrix is a $(\rho^{-1}-1)$-shift of the (partial) unique disjointness matrix, hence for approximations within a factor of ρ, we obtain the lower bound $\mathrm{fc}_+(\mathcal{P}, C, \max) \geq 2^{\frac{n\rho}{8}}$ with Braverman and Moitra [2013], Braun and Pokutta [2013].

See Braun et al. [2014b] for other choices of $\mathcal{G}$, such as e.g., randomly choosing the graphs.

3.1.3 k-juntas via LPs

It is well-known that the level-k Sherali–Adams hierarchy captures all nonnegative k-juntas, i.e., functions $f\colon \{0,1\}^n \to \mathbb{R}_+$ that depend only on k coordinates of the input (see e.g., Chan et al. [2013]) and it can be written as a linear program using $O(n^k)$ inequalities. We will now show that this is essentially optimal for k small.

Example 3.11 (k-juntas). We consider the problem of maximizing nonnegative k-juntas over the n-dimensional hypercube. Let the set of instances $\mathcal{F}$ be the family of all nonnegative k-juntas and let the set of feasible solutions be $\mathcal{S} = \{0,1\}^n$, with $\mathrm{val}_f(s) := \mathrm{val}_f(s)$. We put $C(f) = S(f) = \max_{s\in\mathcal{S}} \mathrm{val}_f(s)$.

As we are interested in a lower bound we will confine ourselves to a specific subfamily of functions $\mathcal{F}' := \{f_a \,|\, a \in \{0,1\}^n, |a| = k\} \subseteq \mathcal{F}$ with

$$f_a(b) := a^\intercal b - 2\binom{a^\intercal b}{2},$$

and hence $C(f_a) = 1$. Clearly $|\mathcal{F}'| = \binom{n}{k}$, so that the nonnegative rank of the slack matrix

$$S_{a,b} := C(f_a) - f_a(b) = 1 - a^\intercal b + 2\binom{a^\intercal b}{2} = (1 - a^\intercal b)^2,$$

with $a, b \in \{0,1\}^n$ and $|a| = k$ is at most $\binom{n}{k}$.

Now for each $f_a \in \mathcal{F}'$ we have that $C(f_a) - f_a(b) = (1 - a^\intercal b)^2 = 1$ if $a \cap b = \emptyset$ and there are 2^{n-k} such choices for b for a given a. Thus the matrix S has $\binom{n}{k}2^{n-k}$ entries 1 arising from disjoint pairs a, b. However in Kaibel and Weltge [2013] it was shown that any nonnegative rank-1 matrix can cover at most 2^n of such pairs. Thus the nonnegative rank of S is at least

$$\frac{\binom{n}{k}2^{n-k}}{2^n} = \frac{\binom{n}{k}}{2^k}.$$

The latter is $\Omega(n^k)$ for k constant and at least $\Omega(n^{k-\alpha})$ for $k = \alpha \log n$ with $\alpha \in \mathbb{N}$ constant and $k > \alpha$. Thus the LP formulation for k-juntas derived from the level-k Sherali-Adams hierarchy is essentially optimal for small k.

4 Affine Reductions for LPs and SDPs

We will now introduce natural reductions between problems, with control on approximation guarantees that translate to the underlying LP and SDP level.

Definition 4.1 (Reductions between problems)**.** Let $\mathcal{P}_1 = (\mathcal{S}_1, \mathcal{F}_1, \mathrm{val})$ and $\mathcal{P}_2 = (\mathcal{S}_2, \mathcal{F}_2, \mathrm{val})$ be maximization problems. Let C_1, S_1 and C_2, S_2 be guarantees for $\mathcal{P}_1$ and $\mathcal{P}_2$ respectively. A *reduction* from $\mathcal{P}_1$ to $\mathcal{P}_2$ respecting these guarantees consist of two maps:

(i) $\beta\colon \mathcal{F}_1^{\mathcal{S}_1} \to \mathrm{cone}\left(\mathcal{F}_2^{\mathcal{S}_2}\right) + \mathbb{R}$ rewriting instances as formal nonnegative combinations: $\beta(f_1) \coloneqq \sum_{f\in\mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f}\cdot f + \mu(f_1)$ with $b_{f_1,f} \geq 0$ for all $f \in \mathcal{F}_2^{\mathcal{S}_2}$; the term $\mu(f_1)$ is called the *affine shift*

(ii) $\gamma\colon \mathcal{S}_1 \to \mathrm{conv}\,(\mathcal{S}_2)$ rewriting solutions as formal convex combination of $\mathcal{S}_2$: $\gamma(s_1) \coloneqq \sum_{s\in\mathcal{S}_2} a_{s_1,s}\cdot s$ with $a_{s_1,s} \geq 0$ for all $s \in \mathcal{S}_2$ and $\sum_{s\in\mathcal{S}_2} a_{s_1,s} = 1$;

subject to

$$\mathrm{val}_{f_1}(s_1) = \sum_{\substack{f\in\mathcal{F}_2^{\mathcal{S}_2}\\ s\in\mathcal{S}_2}} b_{f_1,f}a_{s_1,s}\cdot \mathrm{val}_f(s) + \mu(f_1), \qquad s_1 \in \mathcal{S}_1,\ f_1 \in \mathcal{F}_1^{\mathcal{S}_1}, \tag{8}$$

expressing representation of the objective function of $\mathcal{P}_1$ by that of $\mathcal{P}_2$, and additionally

$$C_1(f_1) \geq \sum_{f\in\mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f}\cdot C_2(f) + \mu(f_1), \qquad f_1 \in \mathcal{F}_1^{\mathcal{S}_1}, \tag{9}$$

ensuring feasibility of the completeness guarantee.

Observe that the role of soundness guarantees of $\mathcal{P}_1$ and $\mathcal{P}_2$ in the definition is to restrict the instances considered: the map β involves only the instances whose optimum value is bounded by these guarantees. One can analogously define reductions involving minimization problems. E.g., for a reduction from a maximization problem $\mathcal{P}_1$ to a minimization problem $\mathcal{P}_2$, the formulas are

$$\begin{aligned}
\beta(f_1) &\coloneqq \mu(f_1) - \sum_{f\in\mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f}\cdot f\\
\mathrm{val}_{f_1}(s_1) &= \mu(f_1) - \sum_{\substack{f\in\mathcal{F}_2^{\mathcal{S}_2}\\ s\in\mathcal{S}_2}} b_{f_1,f}a_{s_1,s}\cdot \mathrm{val}_f(s)\\
C_1(f_1) &\geq \mu(f_1) - \sum_{\mathrm{val}_f\in\mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f}\cdot C(f).
\end{aligned}$$

Note that elements in $\mathcal{S}_2$ are obtained as convex combinations, while elements in $\mathcal{F}_2$ are obtained as nonnegative combinations and a shift. The additional freedom for instances allows scaling and shifting the function values.

In a first step we will verify that a reduction between optimization problems $\mathcal{P}_1$ to $\mathcal{P}_2$ naturally extends to potential LP and SDP formulations.

Proposition 4.2 (Reductions of formulations)**.** *Consider a reduction from an optimization problem $\mathcal{P}_1$ to another one $\mathcal{P}_2$ respecting completion and soundness guarantees C_1, S_1 and C_2, S_2. Then* $\mathrm{fc}_+(\mathcal{P}_1, C_1, S_1) \leq \mathrm{fc}_+(\mathcal{P}_2, C_2, S_2)$ *and* $\mathrm{fc}_\oplus(\mathcal{P}_1, C_1, S_1) \leq \mathrm{fc}_\oplus(\mathcal{P}_2, C_2, S_2)$.

Proof. We will use the notation from Definition 4.1 for the reduction. We only prove the claim for LP formulations and for two maximization problems, as the proof is analogous for SDP formulations and when either or both problems are minimization problems. Let us choose an LP formulation $Ax \leq b$ of $\mathcal{P}_2$ with x^s realizing $s \in \mathcal{S}_2$ and w^f realizing $f \in \mathcal{F}_2^{\mathcal{S}_2}$. For $\mathcal{P}_1$ we shall use the same linear program $Ax \leq b$ with the following realizations y^{s_1} of feasible solutions $s_1 \in \mathcal{S}_1$, and u^{f_1} of instances $f_1 \in \mathcal{F}_1^{\mathcal{S}_1}$, where

$$y^{s_1} := \sum_{s \in \mathcal{S}_2} a_{s_1,s} \cdot x^s, \qquad u^{f_1}(x) := \sum_{f \in \mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f} \cdot w^f(x) + \mu(f_1).$$

As y^{s_1} is a convex combination of the x^s, obviously $Ay^{s_1} \leq b$. The u^{f_1} are clearly affine functions with

$$u^{f_1}(y^{s_1}) = \sum_{f \in \mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f} \cdot w^f \left(\sum_{s \in \mathcal{S}_2} a_{s_1,s} \cdot x^s \right) + \mu(f_1) = \sum_{f \in \mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f} \sum_{s \in \mathcal{S}_2} a_{s_1,s} \cdot w^f(x^s) + \mu(f_1) = \mathrm{val}_{f_1}(s_1)$$

by Eq. (8). Moreover, by Eq. (9).

$$\max_{Ax \leq b} u^{f_1}(x) \leq \sum_{f \in \mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f} \cdot \max_{Ax \leq b} w^f(x) + \mu(f_1) \leq \sum_{f \in \mathcal{F}_2^{\mathcal{S}_2}} b_{f_1,f} \cdot C_2(f) + \mu(f_1) \leq C_1(f_1). \qquad \square$$

Remark 4.3. At the level of matrices, Proposition 4.2 can be equivalently formulated as follows. Whenever $M_1 = R \cdot M_2 \cdot C + t\mathbb{1}$ with M_1, M_2, R, C nonnegative matrices, and t a nonnegative vector, such that $\mathbb{1}C = \mathbb{1}$, then $\mathrm{rank}_{\mathrm{LP}} M_1 \leq \mathrm{rank}_{\mathrm{LP}} M_2$ and $\mathrm{rank}_{\mathrm{SDP}} M_1 \leq \mathrm{rank}_{\mathrm{SDP}} M_2$. Note that given a reduction of $\mathcal{P}_1$ to $\mathcal{P}_2$ with the notation as in Definition 4.1, one chooses M_1 and M_2 to be the slack matrices of $\mathcal{P}_1$ and $\mathcal{P}_2$, respectively, together with matrices R, C and a vector t with the following entries:

$$\begin{aligned} R(f_1, f) &= b_{f_1,f}, \\ C(s, s_1) &= a_{s_1,s}, \\ t(f) &= C_2(f_1) + \mu(f_1) - \sum_{f \in \mathcal{F}_2} b_{f_1,f} \cdot C(f), \end{aligned}$$

all nonnegative, satisfying $M_1 = R \cdot M_2 \cdot C + t\mathbb{1}$. Now we briefly indicate a proof for this alternative formulation of Proposition 4.2.

First, we prove the LP case. Given a size-r LP factorization $M_2 = \sum_{i \in [r]} u_i v_i + \mu\mathbb{1}$ of M_2, one gets the size-r LP factorization $M_1 = \sum_{i \in [r]} Ru_i \cdot v_i C + (R\mu + t)\mathbb{1}$ of M_1.

For the SDP case, let $M_2(i,j) = \mathrm{Tr}[T_i U_j] + \mu(i)$ be an SDP factorization of size r. Then one can construct the following SDP factorization of M_1 of size r:

$$\begin{aligned} M_1(f,s) &= \sum_{i,j} R(f,i) M_2(i,j) C(j,s) + t(f) = \sum_{i,j} R(f,i)(\mathrm{Tr}[T_i U_j] + \mu(i)) C(j,s) + t(f) \\ &= \mathrm{Tr}\left[\underbrace{\left(\sum_i R(f,i) T_i \right)}_{\widehat{T}_f} \cdot \underbrace{\left(\sum_j U_j C(j,s) \right)}_{\widehat{U}_s} \right] + \underbrace{\left(\sum_i R(f,i)\mu(i) + t(f) \right)}_{\widehat{\mu}(f)}, \end{aligned}$$

using $\sum_j C(j,s) = 1$, i.e., $\mathbb{1}C = \mathbb{1}$. Using the nonnegativity of R, C, μ, and t, the $\widehat{T}_f$ and $\widehat{U}_s$ are psd, and the $\widehat{\mu}(f)$ are nonnegative.

For most problems, there is a natural *size* $|f|$ of an instance f, which is a nonnegative number, and guarantees are often proportional to it. In many cases there is no need to consider linear combinations in reductions, leading to a lightweight version of reduction of the form $\beta\colon \mathcal{F}_1 \to \mathcal{F}_2$ and $\gamma\colon \mathcal{S}_1 \to \mathcal{S}_2$. For convenience, the objective function val of $\mathcal{P}_2$ now appears on the left-hand side by rearranging.

Corollary 4.4 (Inapproximability via simple reductions)**.** *Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two optimization problems. Let $\mathcal{P}_1$ the completeness guarantee of $\mathcal{P}_1$ have the form*

$$C_1(f) = \tau_1 |f|, \qquad f \in \mathcal{F}_1$$

proportional to the size $|f|$ of each instance f with $|f| \geq 0$. Furthermore, let $\gamma\colon \mathcal{S}_1 \to \mathcal{S}_2$ and $\beta\colon \mathcal{F}_1^{\mathcal{S}_1} \to \mathcal{F}_2$ be maps satisfying for some constants α, μ

$$\begin{aligned} \operatorname{val}_{\beta(f_1)}[\gamma(s_1)] &= \alpha \operatorname{val}_{f_1}(s_1) + \mu |f_1| \\ |\beta(f_1)| &= (|\alpha| + \mu) \cdot |f_1|, \end{aligned}$$

where $\alpha > 0$ if $\mathcal{P}_2$ is a maximization problem, and $\alpha < 0$ if $\mathcal{P}_2$ is a minimization problem. Let the completeness guarantee C_2 of $\mathcal{P}_2$ be given as

$$C_2(f) \coloneqq (\alpha\tau_1 + \mu) |f| \qquad f \in \mathcal{F}_2.$$

Furthermore let σ_2 be a nonnegative number satisfying for all $f_1 \in \mathcal{F}_1^{\mathcal{S}_1}$

$$\begin{aligned} \max \operatorname{val}_{\beta(f_1)} &\leq \sigma_2 |\beta(f_1)| && \text{if } \mathcal{P}_2 \text{ is a maximization problem} \\ \min \operatorname{val}_{\beta(f_1)} &\geq \sigma_2 |\beta(f_1)| && \text{if } \mathcal{P}_2 \text{ is a minimization problem} \end{aligned}$$

and we set

$$S_2(f) = \sigma_2 |f| \qquad f \in \mathcal{F}_2$$

Then β and γ form a reduction from $\mathcal{P}_1$ with guarantees C_1, S_1 to $\mathcal{P}_2$ with guarantees C_2, S_2. In particular, $\mathcal{P}_2$ is inapproximable within a factor of $\sigma_2/(\alpha\tau_1 + \mu)$ by LP and SDP formulations of size less than $\operatorname{fc}_+(\mathcal{P}_1, C_1, S_1)$ and $\operatorname{fc}_\oplus(\mathcal{P}_1, C_1, S_1)$, respectively.

In many cases also the soundness guarantee of $\mathcal{P}_1$ is proportional to the size of f, i.e.,

$$S_1(f) = \sigma_1 |f|, \qquad f \in \mathcal{F}_1.$$

Then a common choice is $\sigma_2 = \alpha\sigma_1 + \mu$, provided that the reduction is *exact*, i.e.,

$$\operatorname{opt}_{\mathcal{P}_1} \operatorname{val}_{\beta(f_1)} = \alpha \operatorname{opt}_{\mathcal{P}_2} \operatorname{val}_{f_1} + \mu |f_1|,$$

where the operator $\operatorname{opt}_{\mathcal{P}_1}$ is max when $\mathcal{P}_1$ is a maximization problem, and the operator $\operatorname{opt}_{\mathcal{P}_1}$ is min when $\mathcal{P}_1$ is a minimization problem. The operator $\operatorname{opt}_{\mathcal{P}_2}$ is defined similarly for $\mathcal{P}_2$. We shall write $\operatorname{fc}_+(\mathcal{P}, \tau, \sigma)$ for $\operatorname{fc}_+(\mathcal{P}, C, S)$ with $C(f) = \tau |f|$ and $S(f) = \sigma |f|$.

The base problems $\mathcal{P}_1$ from which we reduce will be the CSPs MaxCUT, MaxCUT$_\Delta$ or Max-k-XOR in our examples. For CSPs, the size of an instance, i.e., weighting $(w_1, \ldots, w_m)$ is the total weight $\sum_{i\in[m]} w_i$ of all clauses, For 0/1 weightings representing a subset L of clauses, the size is just the number of elements of L.

The following lower bounds on formulation complexity are implicit in Chan et al. [2013], where similar results are written out explicitly for other constraint satisfaction problems. The problems below constitute our base problems and play the same role as e.g., Max-3-XOR in Håstad's PCP theorem (see Håstad [2001]).

Theorem 4.5. *For every $k \geq 2$ and $\varepsilon > 0$, we have* $\mathrm{fc}_+(\textsf{Max-}k\textsf{-XOR}, 1-\varepsilon, 1/2+\varepsilon)$ *and* $\mathrm{fc}_+(\textsf{MaxCUT}, 1-\varepsilon, 1/2+\varepsilon)$ *are both at least* $n^{\Omega(\frac{\log n}{\log\log n})}$ *for infinitely many n. Moreover, for the bounded degree case we have* $\mathrm{fc}_+(\textsf{MaxCUT}_\Delta, 1-\varepsilon, 1/2+\varepsilon) = n^{\Omega(\frac{\log n}{\log\log n})}$ *for infinitely many n, where Δ is large enough depending on ε.*

Proof. By [Schoenebeck, 2008, Theorem 12], for $k \geq 3$ there are instances L of Max-k-XOR, with at most $\frac{1}{2}+\varepsilon$ of the clauses satisfiable, but with a feasible Lasserre solution of value 1 in round $\Omega(n)$. This means in the language of Chan et al. [2013] that Max-k-XOR is $(1, \frac{1}{2}+\varepsilon)$-inapproximable in the $\Omega(n)$-round Lasserre hierarchy, and therefore also in the Sherali–Adams hierarchy. Therefore by [Chan et al., 2013, Theorem 3.2], Max-k-XOR is also $(1-\varepsilon, \frac{1}{2}+\varepsilon)$-inapproximable by an LP formulation of size $n^{O(\log n/\log\log n)}$ for infinitely many n, which is just a reformulation of the claim for Max-k-XOR.

For MaxCUT and Max-2-XOR the argument is similar, however one uses [Charikar et al., 2009, Proof of Theorem 5.3(I)] to show $(1-\varepsilon, \frac{1}{2}+\varepsilon)$-inapproximability of MaxCUT, and hence of Max-2-XOR. As the construction uses only bounded degree graphs, where the bound Δ depends on ε but not on n, it follows that our argument also works for $\textsf{MaxCUT}_\Delta$. □

Recall from Khot et al. [2007] that under the Unique Games Conjecture, MaxCUT cannot be approximated better than c_{GW} by a polynomial-time algorithm. This motivates the following conjecture, which provides our SDP-hard base problem. For some problems it might be possible to also reduce from Max-3-SAT, which is SDP-hard to approximate within any factor better than $7/8$ Lee et al. [2014a].

Conjecture 4.6 (SDP inapproximability of MaxCUT)**.** *For every $\varepsilon > 0$, and for every constant Δ large enough depending on ε, the formulation complexity* $\mathrm{fc}_\oplus(\textsf{MaxCUT}_\Delta, 1-\varepsilon, c_{GW}+\varepsilon)$ *of* MaxCUT *is super-polynomial.*

Note however that for fixed Δ, there are algorithms achieving an approximation factor of $c_{GW}+\varepsilon$ by Feige et al. [2002], hence in the conjecture Δ should go to infinity as ε tends to 0.

Finally, we remark that by [Karloff, 1999, Lemma 2.9] there are graphs G where the Goemans–Williamson SDP is off by a factor of $c_{GW}+\varepsilon$. For simplicity of calculations, however, we assume for the conjecture that there are also such graphs with SDP optimum $(1-\varepsilon)\,|E(G)|$.

4.1 Facial reductions and formulation complexity

As the notion of formulation complexity does not directly deal with polytopes, there is no direct translation of monotonicity of extension complexity under faces and projections (see Fiorini et al. [2012]). Thus many reductions that have been used in the context of extension complexity and polytopes do not apply, such as e.g., the one from TSP to matching in Yannakakis [1988, 1991]. Often however, the facial reduction underlies a reduction between the problems as defined in Definition 4.1. To exemplify this we provide the underlying reduction from TSP to matching.

Example 4.7 (Maximum weight Hamiltonian cycles (uniform model))*.* We want to find a Hamiltonian cycle with maximum weight in a weighted graph. We consider only nonnegative weights as customary.

Therefore for a fixed n, we choose the *feasible solutions* to be all Hamiltonian cycles C of the complete graph K_n on $[n]$, and the *instances* are weighted subgraphs G of K_n with nonnegative weights. The objective function has the form

$$\mathrm{val}_G(C) \coloneqq \sum_{e \in C \cap E(G)} w_e.$$

We shall consider the exact problem, i.e., with guarantees $C(G) = S(G) = \max \mathrm{val}_G$.

In order to have a finite family of instances, one could restrict the weights to e.g., 1, essentially asking for the maximum number of edges a Hamiltonian cycle can have in common with a given subgraph. For the following reduction, we will use weights $\{1,2\}$ and we adapt Yannakakis's construction to reduce the maximum matching problem on K_{2n} to the maximum weight Hamiltonian cycle problem on K_{4n}. To simplify notation, we identify $[4n]$ with $\{0,1\} \times [2n]$, i.e., the vertices are labelled by pairs (i,j) with i being 0 or 1 and $j \in [2n]$. Given a graph G on $[2n]$, we think of it as being supported on $\{0\} \times [2n]$.

We consider the weighted graph $\widetilde{G}$ with edges and weights:

Edge	Weight	
$\{(0,j),(1,j)\}$	2	$j \in [2n]$
$\{(1,j),(1,k)\}$	1	$j,k \in [2n]$
$\{(0,j),(0,k)\}$	1	$\{j,k\} \in E(G)$

For every perfect matching M on $[2n]$, choose a Hamiltonian cycle C_M containing the edges

(i) $\{(0,j),(1,j)\}$ for $j \in [2n]$,

(ii) $\{(0,j),(0,k)\}$ for $\{j,k\} \in E(G)$,

(iii) n additional edges of the form $\{(1,j),(1,k)\}$ to obtain a Hamiltonian cycle.

Note that

$$\mathrm{val}_{\widetilde{G}}(C_M) = 5n + |M \cap E(G)|\,. \tag{10}$$

We now determine the maximum of $\mathrm{val}_{\widetilde{G}}$ on all Hamiltonian cycles. Therefore let C be an arbitrary Hamiltonian cycle. Let us consider C restricted to $\{0\} \times [2n]$; its components are (possible empty) paths. Let k be the number of components, which are non-empty paths, and contained in G. Obviously, $k \leq \nu(G)$, where $\nu(G)$ is the matching number, as selecting one edge from every such component provides a k-matching of G.

Let l be the number of components containing at least one edge not in G. Note that $k+l \leq n$, because choosing one edge of all these $k+l$ components, we obtain a $(k+l)$-matching on $[2n]$, similarly as in the previous paragraph.

Finally, let m be the number of single vertex components. Therefore C contains exactly $2n-(k+l+m)$ edges on $\{0\} \times [2n]$, of which at least l are not contained in G. Hence the contribution of these edges to the weight $\mathrm{val}_{\widetilde{G}}(C)$ is at most

$$\mathrm{val}_{\widetilde{G}}(C \cap E(\{0\} \times [2n])) \leq 2n-(k+l+m)-l = 2n-k-2l-m. \tag{11}$$

Moreover, the cycle C contains exactly $2n-(k+l+m)$ edges on $\{1\} \times [2n]$ whose contribution to the weight is

$$\mathrm{val}_{\widetilde{G}}(C \cap E(\{1\} \times [2n])) = 2n-(k+l+m). \tag{12}$$

Finally, C contains $2(k+l+m)$ edges between the partitions $\{0\} \times [2n]$ and $\{1\} \times [2n]$, all of which have weight at most 2. In fact, at each of the m single vertex components in $\{0\} \times [2n]$, only one of the edges can be of the form $\{(0,j),(1,j)\}$, the other edge must have weight 0. Therefore the contribution of the edges between the partitions is at most

$$\mathrm{val}_{\widetilde{G}}(C \cap E(\{0\} \times [2n],\{1\} \times [2n]) \leq 2[2(k+l+m)-m] = 4k+4l+2m. \tag{13}$$

Summing up Eqs. (11), (12) and (13), we obtain the following upper bound on the weight of C:

$$\mathrm{val}_{\widetilde{G}}(C) \leq 4n+2k+l \leq 5n+k \leq 5n+\nu(G).$$

Together with Eq. (10), this proves $\max_C \mathrm{val}_{\widetilde{G}}(C) = 5n + \nu(G)$. Thus the $\mathrm{val}_{\widetilde{G}}$ and C_M reduce the maximum matching problem to the maximum Hamiltonian cycle problem with $\beta(G) = \widetilde{G}$, $\gamma(M) = C_M$ and $\mu(G) = 5n$. Hence the LP formulation complexity of the maximum Hamiltonian cycle problem is $2^{\Omega(n)}$ by Proposition 4.2.

5 Inapproximability of VertexCover and IndependentSet

We will now establish inapproximability results for VertexCover and IndependentSet via reduction from MaxCUT, even for bounded degree subgraphs. These two problems are of particular interest, answering a question of Singh [2010] and Chan et al. [2013] as well as a weak version of sparse graph conjecture from Braun et al. [2014a]. Moreover, VertexCover is not of the CSP type, therefore the framework in Chan et al. [2013] does not apply. Using our reduction framework, recently these results have been further improved in Bazzi et al. [2015] to obtain $(2-\varepsilon)$-inapproximability for VertexCover (which is optimal) and inapproximability of IndependentSet within any constant factor.

Formulation complexity depend heavily on how a problem is formulated. For example, the model of IndependentSet used here is motivated by its combinatorial counterpart, and captures standard LPs, like the ones coming from Sherali–Adams hierarchies. In this model, IndependentSet for a given graph G is approximable within a factor of $2\sqrt{n}$ with a polynomial sized LP, see Bazzi et al. [2015]. However, the formulation complexity of *another model of* the maximum independent set problem with an approximation factor $n^{1-\varepsilon}$ is subexponential, rephrasing Fiorini et al. [2012] (see also Braun et al. [2012], Braverman and Moitra [2013], Braun et al. [2014c]), see Section 3.1.2. In this model the instances come from the polytope world, and are actually formal linear combinations of several graphs, and this makes the difference.

The current best PCP bound for bounded degree IndependentSet can be found in Chan [2013]. See also Austrin et al. [2009] for inapproximability results assuming the Unique Games Conjecture.

The minimization problem VertexCover(G) of a graph G asks for a minimum weighted vertex cover of G. We consider the non-uniform model with instances being the induced subgraphs of G.

Definition 5.1 (VertexCover)**.** Given a graph G, the problem VertexCover(G) has all vertex covers S of G as feasible solutions, and instances all induced subgraphs H of G. The problem VertexCover(G) is the minimization problem with its objective function having values $\mathrm{val}_H(S) := |S \cap V(H)|$. The problem VertexCover$(G)_\Delta$ is the restriction of instances to induced subgraphs H, with maximum degree at most Δ.

Note that for every vertex cover S of G, any induced subgraph H has $S \cap V(H)$ as a vertex cover, and all vertex covers of H are of this form. In particular, $\min \mathrm{val}_H$ is the minimum size of a vertex cover of H.

The problem IndependentSet asks for maximum sized independent sets in graphs. As independent sets are exactly the complements of vertex covers, it is natural to use a formulation similar to VertexCover.

Definition 5.2 (IndependentSet)**.** Given a graph G, the maximization problem IndependentSet(G) has all independent sets S of G as feasible solutions, and instances are all induced subgraphs H of G. The objective function is $\mathrm{val}_H(S) := |S \cap V(H)|$. The subproblem IndependentSet$(G)_\Delta$ is the restriction to all induced subgraphs H with maximum degree at most Δ.

For both VertexCover and IndependentSet, we shall use the following conflict graph G for a fixed n, similar to Feige et al. [1991]; we might think of G as a *universal* graph encoding all possible instances. Let the vertices of G be all partial assignments σ of two variables x_i and x_j satisfying the 2-XOR clause $x_i \oplus x_j = 1$. Two vertices σ_1 and σ_2 are connected if and only if the assignments σ_1 and σ_2 are incompatible (i.e., assign different truth values to some variable), see Figure 2 for an illustration. As we are considering problems for optimizing size of vertex sets, it is natural to define the size of an instance, i.e., a subgraph K, as the size of its vertex set $|V(K)|$.

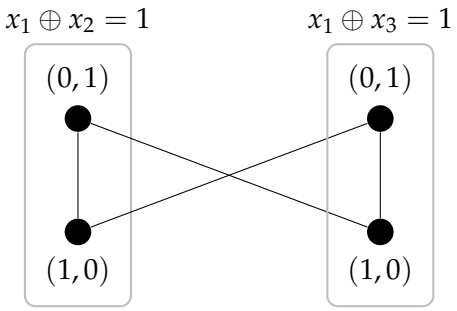

Figure 2: Conflict graph of 2-XOR clauses. We include edges between all conflicted partial assignment to variables.

Theorem 5.3. *For every $\varepsilon > 0$ there is a Δ such that for infinitely many m, there is a graph G with $|V(G)| = m$ such that $\mathrm{fc}_+(\mathsf{VertexCover}(G)_\Delta, 1/2 + \Theta(\varepsilon), 3/4 - \Theta(\varepsilon)) \geq m^{\Omega(\frac{\log m}{\log\log m})}$ showing inapproximability within a factor of $\frac{3}{2} - \Theta(\varepsilon)$, and also $\mathrm{fc}_+(\mathsf{IndependentSet}(G)_\Delta, 1/2 - \Theta(\varepsilon), 1/4 + \Theta(\varepsilon)) \geq m^{\Omega(\frac{\log m}{\log\log m})}$ establishing inapproximability factor $\frac{1}{2} + \Theta(\varepsilon)$. Assuming Conjecture 4.6, we also have $\mathrm{fc}_\oplus(\mathsf{VertexCover}(G)_\Delta, 1/2 + \Theta(\varepsilon), 1 - c_{GW}/2 - \Theta(\varepsilon))$ and $\mathrm{fc}_\oplus(\mathsf{IndependentSet}(G)_\Delta, 1/2 - \Theta(\varepsilon), c_{GW}/2 + \Theta(\varepsilon))$ are superpolynomial, achieving inapproximability factors $2 - c_{GW} - \Theta(\varepsilon)$ and $c_{GW} + \Theta(\varepsilon)$, respectively.*

Proof. We shall use the graph G constructed above, which has $m = 2\binom{n}{2}$ vertices. We reduce MaxCUT_Δ to $\mathsf{VertexCover}(G)_{2\Delta-1}$ using Corollary 4.4 with $\alpha = -1, \mu = 2, \tau_1 = 1 - \varepsilon, \sigma_1 = 1/2 + \varepsilon$. For demonstration purposes, we shall write out the explicit guarantees below. Recall that for a graph K on $[n]$ with maximum degree at most Δ, the guarantees for MaxCUT are $C_{\mathsf{MaxCUT}}(K) = (1 - \varepsilon)\,|E(K)|$ and $S_{\mathsf{MaxCUT}}(K) = (1/2 + \varepsilon)\,|E(K)|$. For $\mathsf{VertexCover}(G)_{2\Delta-1}$, we have the following explicit guarantees:

$$C_{\mathsf{VertexCover}(G)}(H) = (1/2 + \varepsilon/2)\,|V(H)|\,,$$
$$S_{\mathsf{VertexCover}(G)}(H) = (3/4 - \varepsilon/2)\,|V(H)|\,.$$

Let $H(K)$ be the induced subgraph of G on the set $V(H(K)) \coloneqq \{\sigma \mid \{i,j\} \in E(K),\ \operatorname{dom}\sigma = \{x_i, x_j\}\}$ of all partial assignments σ which assign values to variables x_i, x_j corresponding to an edge $\{i,j\}$ of K. In particular, $|V(H(K))| = 2\,|E(K)|$, as there are two partial assignments per each edge $\{i,j\}$.

Note that for every partial assignment σ to x_i and x_j, there are $2\Delta - 1$ partial assignments incompatible with it in $V(H(K))$: exactly one assignment for every edge of K incident to i or j. Thus the maximum degree of $H(K)$ is at most $2\Delta - 1$.

We now define the two maps providing the reduction. Let $\beta(K) \coloneqq H(K)$. For a total assignment s, let $\gamma(s) \coloneqq \{\sigma \mid \sigma \not\subseteq s\}$ be the set of partial assignments incompatible with s; this is clearly a vertex cover.

It remains to show that this is a reduction. For every edge $\{i,j\} \in K$, there are two partial assignments σ to x_i and x_j satisfying $x_i \oplus x_j = 1$. If s satisfies $x_i \oplus x_j = 1$, i.e., $\{i,j\}$ is in the cut induced by s, then exactly one of the σ is compatible with s, otherwise both of the assignments are incompatible. This provides

$$\mathrm{val}^{\mathsf{VertexCover}}_{H(K)}[\gamma(s)] = \left|\{\sigma \mid \sigma \not\subseteq s\}\right| = 2\,|E(K)| - \mathrm{val}^{\mathsf{MaxCUT}}_K(s).$$

To compare optimum values, note that for any vertex cover S of G, the partial assignments $\{\sigma \mid \sigma \notin S\}$ occurring in the complement of S are compatible (as the complement forms a stable set), hence there is a global assignment s of $x_1, \ldots, x_n$ compatible with all of them. In particular, $\gamma(s) \subseteq S$, hence $\mathrm{val}^{\mathsf{VertexCover}}_{H(K)}(S) \geq \mathrm{val}^{\mathsf{MaxCUT}}_K[\gamma(s)]$, so that we obtain

$$\begin{aligned}\min \mathrm{val}^{\mathsf{VertexCover}}_{H(K)} &= \min_s \mathrm{val}^{\mathsf{VertexCover}}_{H(K)}[\gamma(s)] = 2\,|E(K)| - \max \mathrm{val}^{\mathsf{MaxCUT}}_K \\ &\geq 2\,|E(K)| - (1/2 + \varepsilon)\,|E(K)| = (3/4 - \varepsilon/2)\,|V(H(K))| = S_{\mathsf{VertexCover}}(H(K)).\end{aligned}$$

Finally, it is easy to verify that $C_{\mathsf{VertexCover}}(H(K)) = 2\,|E(K)| - C_{\mathsf{MaxCUT}}(K)$. This finishes the proof that β and γ define a reduction to $\mathsf{VertexCover}(G)_\Delta$. Hence by Corollary 4.4 (or Proposition 4.2), using $m = 2\binom{n}{2}$, we have $\mathrm{fc}_+(\mathsf{VertexCover}(G)_{2\Delta-1}, 3/2 - 2\varepsilon, 1 + 2\varepsilon) = n^{\Omega(\frac{\log n}{\log\log n})} = m^{\Omega(\frac{\log m}{\log\log m})}$ for infinitely many n.

For $\mathsf{IndependentSet}$, we apply a similar reduction from MaxCUT_Δ to $\mathsf{IndependentSet}(G)_{2\Delta-1}$. We define $\beta(K) \coloneqq H(K)$ as above and we set $\gamma(s) \coloneqq \{\sigma \mid \sigma \subseteq s\}$ to be the set of partial assignments compatible with the total assignment s, this is clearly an independent set, containing exactly one vertex per satisfied clause. In particular, $\mathrm{val}_{H(K)}[\gamma(s)] = \mathrm{val}_K(s)$. The rest of the argument is analogous to the case of $\mathsf{VertexCover}(G)_\Delta$, and hence omitted. Now the parameters for Corollary 4.4 are $\alpha = 1$, $\mu = 0$, $\tau_1 = 1 - \varepsilon$, and $\sigma_1 = 1/2 + \varepsilon$.

The SDP inapproximability factors follow similarly, by replacing $\frac{1}{2}$ with c_{GW}. □

6 Inapproximability of CSPs

In this section we present example reductions for minimum and maximum constraint satisfaction problems. The results for *binary* Max-CSP*s*, (for CSPs as defined in Definition 2.2) could also be obtained in the LP case from Chan et al. [2013] by combination with the respective Sherali–Adams/Lasserre gap instances Lee [2014]. For simplicity of exposition, we reduce from Max-2-XOR, or sometimes MaxCUT, however by reducing from the subproblem MaxCUT_Δ, we immediately obtain the results for bounded occurrence of literals, with Δ depending on the approximation factor.

6.1 Max-MULTI-k-CUT: a non-binary CSP

The Max-MULTI-k-CUT problem is interesting on its own being a CSP over a non-binary alphabet, thus the framework in Chan et al. [2013] does not readily apply. Note that Max-MULTI-k-CUT is APX-hard, as it contains MaxCUT. The current best PCP inapproximability bound $1 - 1/(34k) + \varepsilon$ is given by Kann et al. [1997].

Here we omit the definition of non-binary CSPs, where the feasible solutions are no longer two-valued assignments, and restrict to Max-MULTI-k-CUT.

Definition 6.1 (Max-MULTI-k-CUT)**.** For fixed positive integers n and k, the problem Max-MULTI-k-CUT has

(i) **feasible solutions:** all partitions of $[n]$ into k sets;

(ii) **instances:** all graphs G with $V(G) \subseteq [n]$.

(iii) **objective function:** for a graph G and a partition p of $[n]$, let $\mathrm{val}_G(p)$ be the number of edges of G whose end points lie in different cells of p.

This differs from a binary CSP only by having a different kind of feasible solutions. Hence it is still natural to define the size of an instance, i.e., graph G, as the number of clauses, i.e, number of edges $|E(G)|$.

Corollary 6.2. *Let $k \geq 3$ be a fixed integer. Then for infinitely many n,*

$$\mathrm{fc}_+(\text{Max-MULTI-}k\text{-CUT}, c(k) + 1 - \varepsilon, c(k) + 1/2 + \varepsilon) \geq n^{\Omega\left(\frac{\log n}{\log\log n}\right)}$$

achieving inapproximability factor $\frac{2c(k)+1}{2c(k)+2} + \Theta(\varepsilon)$, where

$$c(k) \coloneqq \binom{k-2}{2}\left(\binom{k+2}{2} - 3\right) + 2(k-2)\left(\binom{k+2}{2} - 3\right).$$

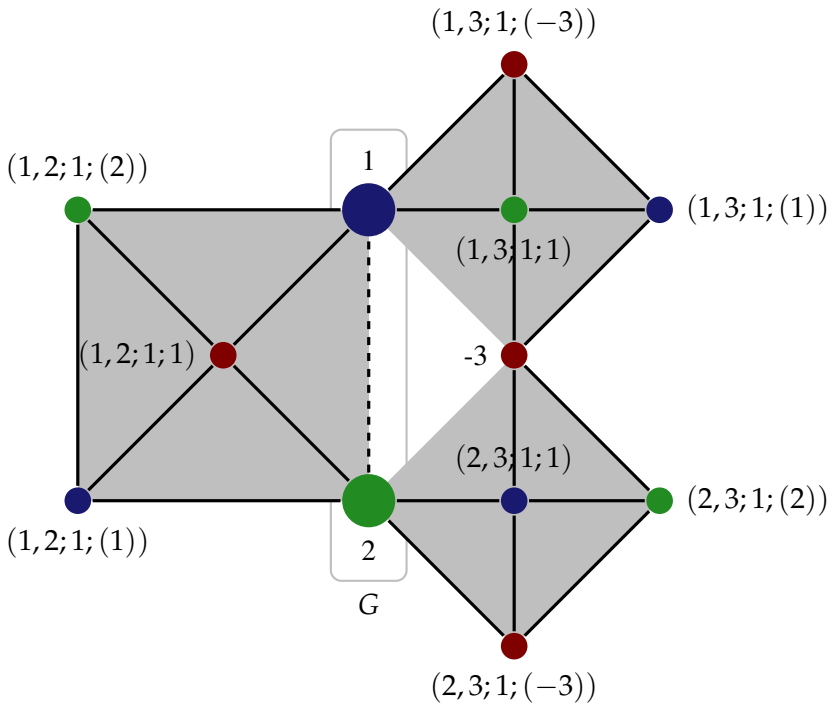

Figure 3: Reduction between MaxCUT and Max-MULTI-k-CUT. Here $k = 3$ and $G = K_2$. The dashed edge denotes the edge of G, which is *not* contained in the reduction. The squares are the copies of almost complete graphs added. The partition is represented by coloring the vertices: blue and green are the original cells on $[2]$, and red is an additional color.

Assuming Conjecture 4.6, we also have that $\mathrm{fc}_{\oplus}(\text{Max-MULTI-}k\text{-CUT}, c(k)+1-\varepsilon, c(k)+c_{GW}+\varepsilon)$ *is superpolynomial, showing inapproximability factor*

$$\frac{c(k)+c_{GW}}{c(k)+1}+\Theta(\varepsilon).$$

Proof. We reduce MaxCUT to Max-MULTI-k-CUT. The reduction is essentially identical to Papadimitriou and Yannakakis [1991], however we have to verify its compatibility with our reduction mechanism. To this end it will suffice to define the reduction maps β and γ.

Given a graph G, we construct a new graph $\beta(G)$ as illustrated in Figure 3. Consider the vertices of G together with $k-2$ new vertices $-k,\dots,-3$. For every pair of vertices i,j we add n_{ij} copies of an almost complete graph on $k+2$ vertices, two of which are $(i),(j)$, as follows.

First, let us determine the number n_{ij} of copies added. For $i,j\in[n]$, we add one copy if i and j are connected in G, and no copies otherwise. For $i,j\in\{-k,\dots,-3\}$ we add $|E(G)|$ many copies. Finally, for $i\in[n]$ and $j\in\{-k,\dots,-3\}$ we add $\deg_G(i)$ many copies.

Now let us describe the copies themselves. Let us fix i,j and let $x\in[n_{ij}]$ be an index of the copy. We add k new vertices $(i,j;x;(i))$, $(i,j;x;(j))$, and $(i,j;x;t)$ for $t=1,\dots,k-2$. We connect every pair of the $k+2$ vertices $i,j,(i,j;x;(i)),(i,j;x;(j)),(i,j;x;t)$ with an edge except the pairs $\{i,(i,j;x,(i))\}$, $\{i,j\}$, and $\{(i,j;x,j),(j)\}$. This is done for all i,j,x, and we let $\beta(G)$ be the graph so obtained.

By construction, $\beta(G)$ has

$$|E(\beta(G))|=\left[\binom{k+2}{2}-3\right]\cdot\sum_{ij}n_{ij}=\left[\binom{k+2}{2}-3\right]\left(1+2(k-2)+\binom{k-2}{2}\right)|E(G)|$$

many edges, and its vertex set $V(\beta(G))$ is contained in the set

$$[n] \cup \{-k, \dots, -3\} \cup \left\{(i,j;x;t) \;\middle|\; -k \le i < j \le 3, x \in [n_{ij}], t \in [k]\right\}$$

having size polynomial in n:

$$m := n + (k-2) + k\left[\binom{k-2}{2} + 2(k-2) + 1\right]\binom{n}{2}.$$

vertices.

We define γ to map every 2-partition p of $[n]$ into a k-partition of $\beta(G)$ by extending it as follows. Let p_1 and p_2 denote the cells of p. Elements of p_1 and p_2 go to the first and second cell of the $\gamma(p)$, respectively. The new vertices $-i$ added for $i = -k, \dots, -3$ go to the i-th cell. Vertices $(i,j;x;(i))$ and $(i,j;x;(j))$ go to the cell of i and j, respectively. For fixed i, j, x the vertices $(i,j;x;t)$ for $t = 1, \dots, k-2$ are put into $k-2$ different cells, which do not contain $(i,j;x;(i))$ or $(i,j;x;(j))$. This is possible as there are k different cells.

By Papadimitriou and Yannakakis [1991],

$$\begin{aligned}\mathrm{val}_{\beta(G)}^{\text{Max-MULTI-}k\text{-CUT}}[\gamma(p)] &= \mathrm{val}_G^{\text{MaxCUT}}(p) + \mu(G),\\ \max \mathrm{val}_{\beta(G)}^{\text{Max-MULTI-}k\text{-CUT}} &= \max \mathrm{val}_G^{\text{MaxCUT}} + \mu(G),\end{aligned}$$

where

$$\mu(G) = |E(\beta(G))| - |E(G)| = \underbrace{\left[\binom{k-2}{2} + 2(k-2)\right]\left[\binom{k+2}{2} - 3\right]}_{c(k)} |E(G)|.$$

Therefore we obtain a reduction from MaxCUT on n vertices to Max-MULTI-k-CUT on m vertices, with m polynomially bounded in n. Combining Corollary 4.4 with Theorem 4.5 in the LP case with parameters $\alpha = 1$, $\mu = c(k)$, $\tau_1 = 1-\varepsilon$, $\sigma_1 = 1/2+\varepsilon$, we obtain that $\mathrm{fc}_+(\text{Max-MULTI-}k\text{-CUT}, c(k)+1-\varepsilon, c(k)+1/2+\varepsilon)$ and is $m^{\Omega(\frac{\log m}{\log\log m})} = n^{\Omega(\frac{\log n}{\log\log n})}$. The SDP case follows analogously from Conjecture 4.6 using $\sigma_1 = c_{GW}/2 + \varepsilon$. □

6.2 Inapproximability of general 2-CSPs

First we consider general CSPs with no restrictions on constraints, for which the exact approximation factor can be easily established. We present the hardness of LP approximation here. The LP with matching factor can be found in Trevisan [1998].

Definition 6.3 (Max-2-CSP and Max-2-CONJSAT)**.** The problem Max-2-CSP is the CSP on variables $x_1, \dots, x_n$ with constraint family $\mathcal{C}_{2CSP}$ consisting of all possible constraints depending on at most two variables. The problem Max-2-CONJSAT is the CSP with constraint family consisting of all possible conjunctions of two literals.

Corollary 6.4. *For every $\varepsilon > 0$ and infinitely many n, we have* $\mathrm{fc}_+(\text{Max-2-CSP}, 1-\varepsilon, 1/2+\varepsilon) \ge n^{\Omega(\frac{\log n}{\log\log n})}$ *achieving inapproximability factor $\frac{1}{2} + \Theta(\varepsilon)$, where n is the number of variables of* Max-2-CSP*. Similarly,* $\mathrm{fc}_+(\text{Max-2-CONJSAT}, 1/2-\varepsilon, 1/4+\varepsilon) \ge n^{\Omega(\frac{\log n}{\log\log n})}$ *establishing inapproximability factor $\frac{1}{2} + \Theta(\varepsilon)$ for infinitely many n. Moreover, in the SDP case, assuming Conjecture 4.6, we have* $\mathrm{fc}_\oplus(\text{Max-2-CSP}, 1-\varepsilon, c_{GW}+\varepsilon)$ *and* $\mathrm{fc}_\oplus(\text{Max-2-CONJSAT}, 1/2-\varepsilon, c_{GW}/2+\varepsilon)$ *are superpolynomial showing inapproximability factor $c_{GW} + \Theta(\varepsilon)$.*

Proof. We identify Max-2-XOR as a subproblem of Max-2-CSP: Every 2-XOR clause is evidently a boolean function of 2 variables. So restricting the instances of Max-2-CSP to 2-XOR clauses with 0/1 weights gives Max-2-XOR. Now with Theorem 4.5, the result follows.

The claim about Max-2-CONJSAT follows via the reduction from Max-2-CSP to Max-2-CONJSAT in Trevisan [1998]. We prefer to reduce from Max-2-XOR instead for easier control over the approximation guarantees. The idea is to write each clause C in disjunctive normal form, and replace C with the set $S(C)$ of conjunctions in its normal form, one conjunction for every assignment satisfying C. In particular, for 2-XOR clauses $S(x_i \oplus x_j = 1) = \{x_i \wedge \neg x_j, \neg x_i \wedge x_j\}$ and $S(x_i \oplus x_j = 0) = \{x_i \wedge x_j, \neg x_i \wedge \neg x_j\}$. Therefore formally, a set of clauses L is mapped to $\beta(L) = \bigcup_{C \in L} S(C)$. Every assignment of variables is mapped to themselves, i.e., γ is the identity. We have $\mathrm{val}_{\beta(L)}(s) = \mathrm{val}_L(s)$ and $|\beta(L)| = 2|L|$. Now the claim follows. □

6.3 Max-2-SAT and Max-3-SAT inapproximability

We now establish an LP-inapproximability factor of $\frac{3}{4} + \varepsilon$ for Max-2-SAT via a direct reduction from MaxCUT and an SDP-inapproximability factor of about $0.93928 + \varepsilon$ assuming Conjecture 4.6. Note that Goemans and Williamson [1994] show the existence of an LP that achieves a factor of $\frac{3}{4}$, so that our estimation is tight in the LP case. Moreover, in Feige and Goemans [1995] it is shown that Max-2-SAT can be approximated with a small SDP within a factor of 0.931 leaving a (conditional) gap of about 0.08.

Obviously, the same factor applies for Max-k-SAT with $k \geq 2$, too. Note that we allow clauses with less than k literals in Max-k-SAT, which is in line with the definition in Schoenebeck [2008] to maintain compatibility. Note that [Lee et al., 2014a, Theorem 1.5] establishes $7/8 + \varepsilon$ inapproximability for Max-3-SAT even in the SDP case.

Definition 6.5 (Max-k-SAT)**.** For fixed $n, k \in \mathbb{N}$, the problem Max-k-SAT is the CSP on the set of variables $\{x_1, \ldots, x_n\}$, where the constraint family $\mathcal{C}$ is the set of all sat clauses which consist of at most k literals.

Corollary 6.6. *For infinitely many n,* $\mathrm{fc}_+(\text{Max-2-SAT}, 1-\varepsilon, 3/4+\varepsilon) \geq n^{\Omega(\frac{\log n}{\log\log n})}$ *and* $\mathrm{fc}_+(\text{Max-3-SAT}, 1-\varepsilon, 3/4+\varepsilon) \geq n^{\Omega(\frac{\log n}{\log\log n})}$ *achieving inapproximability factor $\frac{3}{4} + \Theta(\varepsilon)$, where n is the number of variables. In the case of SDPs, assuming Conjecture 4.6, we have* $\mathrm{fc}_\oplus(\text{Max-2-SAT}, 1-\varepsilon, (1+c_{GW})/2)$ *and* $\mathrm{fc}_\oplus(\text{Max-3-SAT}, 1-\varepsilon, (1+c_{GW})/2)$ *are both superpolynomial establishing inapproximability factor $\frac{1+c_{GW}}{2} + \Theta(\varepsilon) \approx 0.93928 + \Theta(\varepsilon)$.*

Proof. We reduce MaxCUT to Max-2-SAT. For a 2-XOR clause $l = (x_i \oplus x_j = 1)$ with $i, j \in [n]$, we define two auxiliary constraints $C_1(l) = (x_i \vee x_j)$ and $C_2(l) = (\bar{x}_i \vee \bar{x}_j)$. Let $\beta(L) := \{C_1(l), C_2(l) \mid l \in L\}$ for a set of 2-XOR clauses L. We choose γ to be the identity map. Observe that whenever l is satisfied by a partial assignment s then both $C_1(l)$ and $C_2(l)$ are also satisfied by s, otherwise exactly one of $C_1(l)$ and $C_2(l)$ is satisfied. Hence we obtain a reduction from MaxCUT to Max-2-SAT. Together with Theorem 4.5 and Proposition 4.2 the result follows. The statement for Max-3-SAT follows, as Max-2-SAT is a subproblem of Max-3-SAT. □

Max-DICUT inapproximability

Problem Max-DICUT asks for a maximum sized cut in a directed graph G, i.e., partitioning the vertex set $V(G)$ into two parts V_0 and V_1, such that the number of directed edges $(i, j) \in E(G)$ going from V_0 to V_1, i.e., $i \in V_0$ and $j \in V_1$ are maximal. We use a formulation similar to MaxCUT.

Definition 6.7 (Directed Cut)**.** For a fixed $n \in \mathbb{N}$, the problem Max-DICUT is the CSP with constraint family $\mathcal{C}_{DICUT} = \{\neg x_i \wedge x_j \mid i, j \in [n], i \neq j\}$.

We obtain $(1/2+\varepsilon)$-inapproximability via the standard reduction from undirected graphs, by replacing every edge with two, namely, one edge in either direction. The inapproximability factor is tight as the LP in [Trevisan, 1998, Page 84, Eq. (DI)], is $\frac{1}{2}$-approximate for maximum weighted directed cut. In the SDP case we obtain $c_{GW}+\varepsilon$-inapproximability assuming Conjecture 4.6 and in Feige and Goemans [1995] it is shown that Max-DICUT can be approximated with a small SDP within a factor of 0.859 leaving a (conditional) gap of about 0.02.

Corollary 6.8. *For infinitely many n, we have* $\mathrm{fc}_+(\text{Max-DICUT}, 1/2-\varepsilon, 1/4+\varepsilon) \geq n^{\Omega(\frac{\log n}{\log\log n})}$ *achieving inapproximability factor* $1/2+\Theta(\varepsilon)$. *Assuming Conjecture 4.6, in the SDP case* $\mathrm{fc}_{\oplus}(\text{Max-DICUT}, 1/2-\varepsilon, c_{GW}/2+\varepsilon)$ *is superpolynomial establishing inapproximability factor* $c_{GW}+\Theta(\varepsilon)$.

Proof. The proof is analogous to that of Corollaries 6.4 and 6.6, hence we point out only the differences We reduce MaxCUT to Max-DICUT, but now replace every clause $l=(x_i\oplus x_j=1)$ with $C_1(l)=\neg x_i\wedge x_j$ and $C_2(l)=x_i\wedge\neg x_j$. Now observe that whenever $x_i\oplus x_j=1$ is satisfied by a partial assignment s then exactly one of $\neg x_i\wedge x_j$ and $x_i\wedge\neg x_j$ is also satisfied by s. If on the other hand $x_i\oplus x_j$ is not fulfilled by s, then neither $\neg x_i\wedge x_j$ nor $x_i\wedge\neg x_j$ are fulfilled. The remainder of the proof is the same as in Corollary 6.6. □

6.4 Minimum constraint satisfaction

In this section we examine minimum constraint satisfaction problems, a variant of constraint satisfaction problems, where the objective is not to maximize the number of *satisfied* constraints, but to minimize the number of *unsatisfied* constraints. This is equivalent to maximizing the number of satisfied constraints, however, the changed objective function yields different approximation factors due to the change in the magnitude of the optimum value; this is in analogy to the algorithmic world. We consider only Min-2-CNFDeletion and MinUnCUT from Agarwal et al. [2005], which are complete in their class in the algorithmic hierarchy; our technique applies to many more problems in Papadimitriou and Yannakakis [1991]. The problem Min-2-CNFDeletion is of particular interest here, as it is considered to be the hardest minimum CSP with nontrivial approximation guarantees (see Agarwal et al. [2005]). We start with the general definition of minimum CSPs.

Definition 6.9. The *minimum Constraint Satisfaction Problem* on variables $x_1,\ldots,x_n$ with constraint family $\mathcal{C}=\{C_1,\ldots,C_m\}$ is the minimization problem with

(i) **feasible solutions** all 0/1 assignments to $x_1,\ldots,x_n$;

(ii) **instances** all nonnegative weightings $w_1,\ldots,w_m$ of the constraints $C_1,\ldots,C_m$;

(iii) **objective functions** weighted sum of negated constraints, i.e. $\mathrm{val}_{w_1,\ldots,w_m}(x_1,\ldots,x_n)=\sum_i w_i[1-C_i(x_1,\ldots,x_n)]$.

The goal is to minimize the objective function, i.e., the weight of unsatisfied constraints.

As mentioned above, we consider two examples.

Example 6.10 (Minimum CSPs). The problem Min-2-CNFDeletion is the minimum CSP with constraint family consisting of all disjunction of two literals, as in Max-2-SAT. The problem MinUnCUT is the minimum CSP with constraint family consisting of all equations $x_i\oplus x_j=b$ with $b\in\{0,1\}$, as in Max-2-XOR.

We are ready to prove LP inapproximability bounds for these problems. Instead of the reductions in Chlebík and Chlebíková [2004], we use direct, simpler reductions from MaxCUT and here we provide reductions for general weights. Note that the current best known algorithmic inapproximability for Min-2-CNFDeletion is $8\sqrt{5}-15-\varepsilon\approx 2.88854-\varepsilon$ by Chlebík and Chlebíková [2004]. Assuming the Unique Games Conjecture, Chawla et al. [2006] establish that Min-2-CNFDeletion cannot be approximated within any constant

factor and our LP inapproximability factor coincides with this one. The problem MinUnCUT is known to be SNP-hard (see Papadimitriou and Yannakakis [1991]) however the authors are not aware of the strongest known factor. We refer the reader to Khanna et al. [2001] for a classification of all minimum CSPs.

Theorem 6.11. *For every $\varepsilon > 0$ and infinitely many n, we have* $\mathrm{fc}_+(\textsf{Min-2-CNFDeletion}, \varepsilon, 1/4 - \varepsilon) = n^{\Omega(\frac{\log n}{\log\log n})}$ *and* $\mathrm{fc}_+(\textsf{MinUnCUT}, \varepsilon, 1/2 - \varepsilon) = n^{\Omega(\frac{\log n}{\log\log n})}$ *establishing inapproximability within any constant factor, where n is the number of variables. Assuming Conjecture 4.6, even in the SDP case we have that* $\mathrm{fc}_\oplus(\textsf{Min-2-CNFDeletion}, \varepsilon, (1 - c_{GW})/2 - \varepsilon)$ *and* $\mathrm{fc}_\oplus(\textsf{MinUnCUT}, \varepsilon, 1 - c_{GW} - \varepsilon)$ *are superpolynomial showing inapproximability within any constant factor.*

Proof. We reduce from MaxCUT to Min-2-CNFDeletion similar to the previous reductions: assignments are mapped to themselves, i.e., γ is the identity. Under β every clause C_ℓ is replaced with two disjunctive clauses $C_\ell(1)$ and $C_\ell(2)$, both inheriting the weight w_ℓ of C_ℓ, i.e., $\mathrm{val}^{\textsf{Min-2-CNFDeletion}}_{\beta(w_1,\ldots,w_m)}(x_1,\ldots,x_n) = \sum_\ell w_\ell[(1 - C_\ell(1)) + (1 - C_\ell(2))]$.

For $C_\ell = (x_i \oplus x_j = 1)$, we let $C_\ell(1) := x_i \vee \neg x_j$ and $C_\ell(2) := \neg x_i \vee x_j$. Note that if C_ℓ is unsatisfied, then both of $C_\ell(1)$ and $C_\ell(2)$ are satisfied, and if C_ℓ is satisfied, then exactly one of $C_\ell(1)$ and $C_\ell(2)$ is satisfied. Therefore $\mathrm{val}^{\textsf{Min-2-CNFDeletion}}_{\beta(w_1,\ldots,w_m)} = \sum_\ell w_\ell - \mathrm{val}^{\textsf{MaxCUT}}_{w_1,\ldots,w_m} w_\ell$. This provides the desired lower bound.

For MinUnCUT the reduction is similar but simpler, as we replace every clause C with itself. □

7 From matrix approximation to problem approximations

We will now explain how a nonnegative matrix with small nonnegative rank (or semidefinite rank) that is close to a slack matrix of a problem $\mathcal{P}$ of interest can be rounded to an actual slack matrix with a moderate increase in nonnegative rank (or semidefinite rank) and error. This argument is implicitly contained in Rothvoß [2011], Briët et al. [2014] for the linear and semidefinite case respectively. In some sense we might want to think of this approach as an interpolation between a slack matrix (which corresponds to $\mathcal{P}$) and a close-by matrix of low nonnegative rank (or semidefinite rank) that does not correspond to any optimization problem. The result is a low nonnegative rank (or semidefinite rank) approximation of $\mathcal{P}$ with small error.

We will need the following simple lemma. Recall that the exterior algebra of a vector space V is the $\mathbb{R}$-algebra generated by V subject to the relations $v^2 = 0$ for all $v \in V$. As is customary, the product in this algebra is denoted by $\wedge$. The subspace of homogeneous degree-k elements (i.e., linear combination of elements of the form $v_1 \wedge \cdots \wedge v_k$ with $v_1, \ldots, v_k \in V$) is denoted by $\bigwedge^k V$. Recall that for $k = \dim V$, the space $\bigwedge^k V$ is one dimensional and is generated by $v_1 \wedge \cdots \wedge v_k$ for any basis $v_1, \ldots, v_k$ of V.

Lemma 7.1. *Let $M \in \mathbb{R}^{m\times n}$ be a real matrix of rank r. Then there are column vectors $a_1, \ldots, a_r \in \mathbb{R}^m$ and row vectors $b_1, \ldots, b_r \in \mathbb{R}^n$ with $M = \sum_{i\in[r]} a_i b_i$. Moreover, $\|a_i\|_\infty \leq 1$ and $\|b_i\|_\infty \leq \|M\|_\infty$ for all $1 \leq i \leq r$.*

Proof. Consider the r dimensional vector space V spanned by all the rows $M_1, \ldots, M_m$ of M and identify the one dimensional exterior product $\bigwedge^r V$ with $\mathbb{R}$. Now choose r rows $M_{i_1}, \ldots, M_{i_r}$ for which $M_{i_1} \wedge \cdots \wedge M_{i_r}$ is the largest in absolute value in $\mathbb{R}$. As the M_i together span V it follows that the largest value is non-zero. Hence $M_{i_1}, \ldots, M_{i_r}$ form a basis of V. Therefore any row M_k can be uniquely written as a linear combination of the basis elements:

$$M_k = \sum_{j\in[r]} a_{k,j} M_{i_j}. \tag{14}$$

Fixing $j \in [r]$ and taking exterior products with the M_{i_l} where $l \neq j$ and both side we obtain

$$M_{i_1} \wedge \ldots \wedge M_{i_{j-1}} \wedge M_k \wedge M_{i_{j+1}} \wedge \ldots \wedge M_{i_r} = a_{k,j} \cdot M_{i_1} \wedge \ldots \wedge M_{i_r},$$

using the vanishing property of the exterior product. By maximality of $M_{i_1} \wedge \ldots \wedge M_{i_r}$, it follows that $|a_{k,j}| \leq 1$. We choose $a_k := \begin{bmatrix} a_{k,1} \\ \cdots \\ a_{k,r} \end{bmatrix}$, and thus we have $\|a_k\|_\infty \leq 1$. Moreover choose $b_j := M_{i_j}$, so that $\|b_j\|_\infty \leq \|M\|_\infty$ holds. Finally, Eq. (14) can be rewritten to $M = \sum_{j\in[r]} a_j b_j$, finishing the proof. □

For a vector a we can decompose it into its positive and negative part so that $a = a^+ - a^-$ with $a^+ a^- = 0$. Let $|a|$ denote the vector obtained from a by replacing every entry with its absolute value. Note that a^+, a^-, and $|a|$ are nonnegative vectors and $|a| = a^+ + a^-$. Furthermore their ℓ_∞-norm is at most $\|a\|_\infty$.

Theorem 7.2. *Let $\mathcal{P}$ be an optimization problem with (C, S)-approximate slack matrix M and let $\widetilde{M}$ be a nonnegative matrix. Then for the adjusted guarantee C' for $\mathcal{P}$ defined as*

$$C'(f) := C(f) + (\operatorname{rank} M + \operatorname{rank} \widetilde{M})\|\widetilde{M} - M\|_\infty \qquad \text{if } \mathcal{P} \text{ is a maximization problem, and}$$
$$C'(f) := C(f) - (\operatorname{rank} M + \operatorname{rank} \widetilde{M})\|\widetilde{M} - M\|_\infty \qquad \text{if } \mathcal{P} \text{ is a minimization problem,}$$

we have

$$\mathrm{fc}_+(\mathcal{P}, C', S) \leq \operatorname{rank}_{\mathrm{LP}} \widetilde{M} + 2(\operatorname{rank} M + \operatorname{rank} \widetilde{M}) \qquad \text{and}$$
$$\mathrm{fc}_\oplus(\mathcal{P}, C', S) \leq \operatorname{rank}_{\mathrm{SDP}} \widetilde{M} + 2(\operatorname{rank} M + \operatorname{rank} \widetilde{M}).$$

Proof. We prove the statement for maximization problems; the minimization case follows similarly. The proof is based on the vector identity

$$\sum_{i\in[k]} |a_i| b - \sum_{i\in[k]} a_i b_i = \sum_{i\in[k]} a_i^+ (b - b_i) + \sum_{i\in[k]} a_i^- (b + b_i). \tag{15}$$

In our setting, the a_i, b_i with $i \in [k]$ will arise from the (not necessarily nonnegative) factorization of $\widetilde{M} - M$, obtained by applying Lemma 7.1, i.e., we have

$$\widetilde{M} - M = \sum_{i\in[k]} a_i b_i, \tag{16}$$

where $\|a_i\|_\infty \leq 1$ and $\|b_i\|_\infty \leq \|M\|_\infty$ for $i \in [k]$ with $k \leq \operatorname{rank}(\widetilde{M} - M) \leq \operatorname{rank} M + \operatorname{rank} \widetilde{M}$. Furthermore, define $b := \|\widetilde{M} - M\|_\infty \mathbb{1}$ to be the row vector with all entries equal to $\|\widetilde{M} - M\|_\infty$.

Substituting these values into (15), using Eq. (16) we obtain after rearranging

$$N := \sum_{i\in[k]} |a_i| b + M = \widetilde{M} + \sum_{i\in[k]} a_i^+ (b - b_i) + \sum_{i\in[k]} a_i^- (b + b_i),$$

so that we can conclude

$$\operatorname{rank}_{\mathrm{LP}} N \leq \operatorname{rank}_{\mathrm{LP}} \widetilde{M} + 2k \leq \operatorname{rank}_{\mathrm{LP}} \widetilde{M} + 2(\operatorname{rank} M + \operatorname{rank} \widetilde{M}), \tag{17}$$

and similarly

$$\operatorname{rank}_{\mathrm{SDP}} N \leq \operatorname{rank}_{\mathrm{SDP}} \widetilde{M} + 2k \leq \operatorname{rank}_{\mathrm{SDP}} \widetilde{M} + 2(\operatorname{rank} M + \operatorname{rank} \widetilde{M}). \tag{18}$$

It remains to relate N to the (C', S)-approximate slack matrix of $\mathcal{P}$. By definition, the entries of N are

$$N(f, s) = \underbrace{\sum_{i\in[k]} |a_i(f)| \cdot \|\widetilde{M} - M\|_\infty + C(f)}_{:= f^* \leq C'(f)} - \mathrm{val}_f(s),$$

where $a_i(f)$ is the f-entry of a_i. Furthermore, as $\|a_i\|_\infty \leq 1$ and $k \leq \operatorname{rank} M + \operatorname{rank} \widetilde{M}$, we have $f^* \leq C'(f)$. Thus the (C', S)-approximate slack matrix M' of $\mathcal{P}$ looks like

$$M'(f,s) = N(f,s) + (C'(f) - f^*),$$

and as $f^* \leq C'(f)$, we have $\operatorname{rank}_{\mathrm{LP}} M' \leq \operatorname{rank}_{\mathrm{LP}} N$ and $\operatorname{rank}_{\mathrm{SDP}} M' \leq \operatorname{rank}_{\mathrm{SDP}} N$, establishing the claimed complexity bounds due to (17) in the LP case and (18) in the SDP case. □

A possible application of Theorem 7.2 is to 'thin-out' a given factorization of a slack matrix to obtain an approximation with low nonnegative rank. The idea is that if a nonnegative matrix factorization contains a large number of factors that contribute only very little to each of the entries, then we can simply drop those factors, significantly reduce the nonnegative rank, and obtain a very good approximation of the original optimization problem. Theorem 7.2 is then used to turn the approximation of the matrix into an approximation of the original problem of interest.

Also, it is possible to obtain low rank approximations of combinatorial problems sampling rank-1 factors proportional to their ℓ_1-weight as done in the context of information-theoretic approximations in Braun et al. [2013]. However, the obtained approximations tend to be too weak to be of interest.

8 Final Remarks

We conclude with the following remarks:

(i) Our reduction mechanism works for both LPs and SDPs alike. Unfortunately, there are few inapproximability results for SDPs (such as the one in Lee et al. [2014a]), so that our SDP inapproximability factors are conditional on Conjecture 4.6, which serves as a blackbox here.

(ii) We expect that via an appropriate reduction one can establish LP inapproximability of the TSP, however the current reductions, e.g., in [Engebretsen and Karpinski, 2001, §2] and [Karpinski et al., 2013, §6], cannot be used as they translate feasible solutions depending on the objective functions.

(iii) Our reduction could also establish high formulation complexity for problems in P by reducing from the matching problem.

Acknowledgements

Research reported in this paper was partially supported by NSF grant CMMI-1300144 and NSF CAREER grant CMMI-1452463. The authors would like to thank James Lee for the helpful discussions regarding max-CSPs. We are indebted to Siu On Chan for some of the PCP inapproximability bounds as well as Santosh Vempala for the helpful discussions.

References

A. Agarwal, M. Charikar, K. Makarychev, and Y. Makarychev. $O(\sqrt{\log n})$ approximation algorithms for Min UnCut, Min 2CNF Deletion, and directed cut problems. In *Proceedings of STOC*, pages 573–581. ACM, 2005.

R. K. Ahuja and J. B. Orlin. Inverse optimization. *Operations Research*, 49(5):771–783, 2001.

P. Austrin, S. Khot, and M. Safra. Inapproximability of vertex cover and independent set in bounded degree graphs. In *Computational Complexity, CCC*, pages 74–80. IEEE, 2009.

D. Avis and H. R. Tiwary. On the extension complexity of combinatorial polytopes. *ArXiv:1302.2340*, Feb. 2013.

A. Bazzi, S. Fiorini, S. Pokutta, and O. Svensson. No small linear program approximates vertex cover within a factor $2 - \varepsilon$. *arXiv:1503.00753v1 [cs.CC]*, 2015.

G. Braun and S. Pokutta. Common information and unique disjointness. In *Proceedings of FOCS*, pages 688–697, 2013. doi: 10.1109/FOCS.2013.79. `http://eccc.hpi-web.de/report/2013/056/`.

G. Braun and S. Pokutta. The matching polytope does not admit fully-polynomial size relaxation schemes. *to appear in Processings of SODA / preprint available at* `http://arxiv.org/abs/1403.6710`, 2015.

G. Braun, S. Fiorini, S. Pokutta, and D. Steurer. Approximation limits of linear programs (beyond hierarchies). In *Proceedings of FOCS*, pages 480–489, 2012. ISBN 978-1-4673-4383-1. doi: 10.1109/FOCS.2012.10.

G. Braun, R. Jain, T. Lee, and S. Pokutta. Information-theoretic approximations of the nonnegative rank. *submitted / preprint available at ECCC* `http://eccc.hpi-web.de/report/2013/158/`., 2013.

G. Braun, S. Fiorini, and S. Pokutta. Average case polyhedral complexity of the maximum stable set problem. *Proceedings of RANDOM / arXiv:1311.4001*, 2014a.

G. Braun, S. Fiorini, and S. Pokutta. Average case polyhedral complexity of the maximum stable set problem. *submitted*, 2014b.

G. Braun, S. Fiorini, S. Pokutta, and D. Steurer. Approximation limits of linear programs (beyond hierarchies). *Mathematics of Operations Research*, 2014c. doi: 10.1287/moor.2014.0694. DOI: 10.1287/moor.2014.0694.

G. Braun, J. Brown-Cohen, A. Huq, S. Pokutta, P. Raghavendra, A. Roy, B. Weitz, and D. Zink. The matching problem has no small symmetric SDP. *arXiv:1504.00703 [cs.CC]*, April 2015a.

G. Braun, J. Ostrowski, S. Pokutta, and D. Zink. LP hardness of approximation for (approximate) graph isomorphisms. submitted, 2015b.

G. Braun, S. Pokutta, and A. Roy. Strong reductions for extended formulations. *arXiv preprint arXiv:1512.04932*, 2015c.

M. Braverman and A. Moitra. An information complexity approach to extended formulations. In *Proceedings of STOC*, pages 161–170, 2013.

J. Briët, D. Dadush, and S. Pokutta. On the existence of 0/1 polytopes with high semidefinite extension complexity. *to appear in Mathematical Programming B*, 2014.

S. Chan, J. Lee, P. Raghavendra, and D. Steurer. Approximate constraint satisfaction requires large LP relaxations. In *Proceedings of FOCS*, pages 350–359, 2013. doi: 10.1109/FOCS.2013.45.

S. O. Chan. Approximation resistance from pairwise independent subgroups. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC, pages 447–456, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2029-0. doi: 10.1145/2488608.2488665.

M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for Sherali-Adams relaxations. In *Proceedings of STOC*, pages 283–292, 2009.

S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *computational complexity*, 15(2):94–114, 2006.

M. Chlebík and J. Chlebíková. On approximation hardness of the minimum 2sat-deletion problem. In *Mathematical Foundations of Computer Science 2004*, pages 263–273. Springer, 2004.

J. Edmonds. Maximum matching and a polyhedron with 0, 1 vertices. *Journal of Research National Bureau of Standards*, 69B:125–130, 1965.

L. Engebretsen and M. Karpinski. Approximation hardness of TSP with bounded metrics. In *Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, pages 201–212. Springer Berlin Heidelberg, 2001. ISBN 978-3-540-42287-7 (print) and 978-3-540-48224-6 (online). doi: 10.1007/3-540-48224-5_17. URL `http://theory.cs.uni-bonn.de/ftp/reports/cs-reports/2000/85220-CS.pdf`.

U. Feige and M. Goemans. Approximating the value of two power proof systems, with applications to max 2sat and max dicut. In *Theory of Computing and Systems, 1995. Proceedings., Third Israel Symposium on the*, pages 182–189. IEEE, 1995.

U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating clique is almost NP-complete. In *Proceedings of FOCS*, pages 2–12. IEEE Comput. Soc. Press., October 1991. ISBN 0-8186-2445-0. doi: 10.1109/SFCS.1991.185341.

U. Feige, M. Karpinski, and M. Langberg. Improved approximation of Max-Cut on graphs of bounded degree. *Journal of Algorithms*, 43(2):201–219, 2002. ISSN 0196-6774. doi: 10.1016/S0196-6774(02)00005-6. URL `http://www.paradise.caltech.edu/~mikel/papers/deg_cut.ps`.

S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds. *Proceedings of STOC*, 2012.

M. X. Goemans and D. P. Williamson. New $3/4$-approximation algorithms for the maximum satisfiability problem. *SIAM J. Disc. Math.*, 7(4):656–666, Nov. 1994.

M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42:1115–1145, 1995. doi: 10.1145/227683.227684.

J. Gouveia, P. A. Parrilo, and R. Thomas. Lifts of convex sets and cone factorizations. *Mathematics of Operations Research*, 38(2):248–264, 2013.

J. Håstad. Clique is hard to approximate within $1 - \varepsilon$. *Acta Mathematica*, 182(1):105–142, 1999.

J. Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

V. Kaibel and S. Weltge. A short proof that the extension complexity of the correlation polytope grows exponentially. *ArXiv e-prints*, July 2013.

V. Kaibel, M. Walter, and S. Weltge. LP formulations that induce extensions. private communication, 2013.

V. Kann, S. Khanna, J. Lagergren, and A. Panconesi. On the hardness of approximating Max-k-CUT and its dual. *Chicago Journal of Theoretical Computer Science*, 1997(2):1–18, June 1997.

H. Karloff. How good is the Goemans–Williamson MAX CUT algorithm? *SIAM Journal on Computing*, 29(1):336–350, 1999. doi: 10.1137/S0097539797321481.

M. Karpinski, M. Lampis, and R. Schmied. New inapproximability bounds for TSP. In *Algorithms and Computation*, volume 8283 of *Lecture Notes in Computer Science*, pages 568–578. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-45029-7 (print) and 978-3-642-45030-3 (online). doi: 10.1007/978-3-642-45030-3_53.

S. Khanna, M. Sudan, L. Trevisan, and D. P. Williamson. The approximability of constraint satisfaction problems. *SIAM Journal on Computing*, 30(6):1863–1920, 2001.

S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. doi: 10.1137/S0097539705447372.

J. Lee, P. Raghavendra, and D. Steurer. Lower bounds on the size of semidefinite programming relaxations. *STOC 2015, arXiv:1411.6317*, 2014a.

J. R. Lee. personal communication, Nov. 2014.

J. R. Lee, P. Raghavendra, D. Steurer, and N. Tan. On the power of symmetric LP and SDP relaxations. In *Proceedings of the 2014 IEEE 29th Conference on Computational Complexity*, pages 13–21. IEEE Computer Society, 2014b.

C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of computer and system sciences*, 43(3):425–440, 1991.

K. Pashkovich. *Extended Formulations for Combinatorial Polytopes*. PhD thesis, Magdeburg Universität, 2012.

S. Pokutta and M. Van Vyve. A note on the extension complexity of the knapsack polytope. *Operations Research Letters*, 41:347–350, 2013.

T. Rothvoß. Some 0/1 polytopes need exponential size extended formulations, 2011. arXiv:1105.0036.

T. Rothvoß. The matching polytope has exponential extension complexity. *Proceedings of STOC*, 2014.

G. Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *Proceedings of FOCS*, pages 593–602, Oct. 2008. doi: 10.1109/FOCS.2008.74.

A. Schrijver. *Theory of Linear Programming*. Wiley-Interscience, 1986.

M. Singh. Bellairs workshop on approximation algorithms. Open problem session #1, 2010.

L. Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, 1998. ISSN 0178-4617 (print) and 1432-0541 (online). doi: 10.1007/PL00009209.

L. Trevisan. Inapproximability of combinatorial optimization problems. *The Computing Research Repository*, 2004.

L. Trevisan, G. B. Sorkin, M. Sudan, and D. P. Williamson. Gadgets, approximation, and linear programming. *SIAM Journal on Computing*, 29(6):2074–2097, 2000.

M. Tulsiani. Csp gaps and reductions in the lasserre hierarchy. In *Proceedings of STOC*, pages 303–312. ACM, 2009.

M. Yannakakis. Expressing combinatorial optimization problems by linear programs (extended abstract). In *Proceedings of STOC*, pages 223–228, 1988.

M. Yannakakis. Expressing combinatorial optimization problems by linear programs. *J. Comput. System Sci.*, 43(3):441–466, 1991. doi: 10.1016/0022-0000(91)90024-Y.